Methods for the Expansion of Additive Manufacturing Process Space and the Development of In-Situ Process Monitoring Methodologies

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Mechanical Engineering

Luke Robson Scime

B.S., Mechanical Engineering, The University of Florida M.S., Mechanical Engineering, Carnegie Mellon University

> Carnegie Mellon University Pittsburgh, PA

> > May, 2018

© Luke Robson Scime, 2018

All Rights Reserved

Dedication

To my parents, Earl and Joy Lynn, and my brother Ethan To my dear friend and mentor, Phillip Mark Tucker

Rok tree nol! Meg tal! Mekta satak Oz!

Acknowledgments

This dissertation was by no means a solitary endeavor. Indeed, it was only possible through the kindness and support of others, both professionally and personally. I would first like to thank my advisor and committee chair Prof. Jack Beuth for his advice and continued faith in my abilities. I would also like to thank my other committee members, Prof. Elizabeth Holm, Prof. Levent Burak Kara, Prof. Anthony Rollett, and Dr. Sandra DeVincent Wolf for graciously offering their time and wisdom throughout this entire journey.

My time at Carnegie Mellon University (CMU) has been funded by a variety of sources including CMU's Department of Mechanical Engineering, Mr. Richard Fieler¹, the Research for Advanced Manufacturing in Pennsylvania program², the National Additive Manufacturing Innovation Institute (America Makes)³, CMU's NextManufacturing Center⁴, and CMU's Manufacturing Futures Initiative⁵. Much of the data analyzed in this thesis was incidentally collected during builds sponsored by a wide range of entities; each such sponsor is cited in the body of the text where appropriate. I have also received prior approval for the use of Figure 1.1 from 3D Hubs' copyright department.

Over the last few years I have been extremely fortunate to work with many fantastic collaborators. Much of the presented AlSi10Mg work was done in parallel with related work performed by Dr. Sneha Prabha Narra (CMU) while the experiments themselves were run by

¹ Carnegie Institute of Technology alumni, class of 1956.

² Prime award number FA8650-12-2-7230, sub-award number 543105-78001.

³ Project number 4028.001.

⁴ Internal grant number 39876.2.1011864 and sponsored by General Electric.

⁵ Internal grant number 062900.005.105.100020.01.

Robert Speer, Martin Marinack, and John Siemon at the Arconic Technical Center. Dr. Ming Tang (CMU) assisted with the development of an appropriate AlSi10Mg sample preparation procedure and Recep Önler (CMU) provided initial training on the Alicona Infinite-Focus optical microscope. Characterization of the discussed non-standard powder systems was performed by Hengfeng Gu and Prof. Ola Harrysson of North Carolina State University as well as Ross Cunningham, Prof. Anthony Rollett, Dr. Prathamesh Desai, and Prof. Fred Higgs of CMU. Gregory Le Mon (CMU) performed initial work relating to the design of the test artifacts used to study the non-standard powders. Many of the builds analyzed in thesis were performed by either Dr. Colt Montgomery or Todd Baer of CMU. Brian Fisher (CMU) developed and characterized the high speed camera system used to monitor in-situ melt pool morphologies. Christopher Kantzos (CMU) provided great insight into surface roughness metrics while Dr. Zhou Chen and Prof. Mahadev Satyanarayanan of CMU are working to extend the presented melt pool morphology analysis into the real-time domain. I would also like to thank Prof. Deva Ramanan (CMU) for teaching a course (16–720) on Computer Vision which provided the foundation for some of this work.

I am greatly indebted to the support provided by my "Beuth Lab" colleagues including Debomita Basu, Daniel Christiansen, Daniel Clymer, Daming Ding, Dustin Ferzacca, Brian Fisher, Dr. Zachary Francis, Dr. Jason Fox, Nicholas Jones, Gregory Le Mon, Edgar Jimenez Mendoza, Michael Messersmith, Madeline Miller, Dr. Colt Montgomery, Dr. Sneha Prabha Narra, Shrivani Pandiya, Samuel Passel, Rahi Patel, and José Veintimilla. I would especially like to thank Nick, Edgar, Rahi and Sneha for their help in editing this thesis as well as both Brian and Sneha for their always-invaluable insight. None of my work would have been possible without

v

immeasurable assistance from countless staff at CMU. I would like to thank Edward Wojciechowski and Jim Dillinger for their reliable sass, their wit, and their help machining parts even when it was at the last minute. Todd Baer's dedication and patience were absolutely critical to the smooth operation of the lab. Bobbi Kostyak and Ginny Barry routinely dealt with my last-minute purchasing deadlines and recalcitrant suppliers while John Fulmer and Larry Hayhurst machined an unending stream of build plates. Bill Pingitore helped with etching samples and general metrology techniques. Jeffrey Harris, Michael Patrick, and Tara Balonic succeeded in preventing me from setting the lab on fire. Hannah Diorio-Toth worked creatively to share our group's research with the broader world. Stephen Mrdjenovich, Michael Kelly, Mark Rayburn, Kevin Shaffer, Cliff Humes, Bob Kappas, Tim Kunz, and Bob Oddis all worked tirelessly to make our new lab a reality. During my time at CMU, several faculty members took a special interest in my career, and for that I am particularly grateful to Prof. Reeja Jayan, Prof. Levent Burak Kara, Prof. Allen Robinson, and Prof. Anthony Rollett.

As many of you know, graduate school is more a marathon than a sprint; fortunately I have been blessed to have many friends who have held out water and snacks for me along the way. Sandra, thank you for your guidance and mentoring, you went above and beyond anything I could have ever expected. Sneha, thank you for your unwavering support and counsel. Brian, thank you for your insightful advice, both solicited and otherwise. Lea, thank you for our lunches, your Star Trek memes, and your literal snacks. David, thank you for our late night Facebook messages and your sometimes impenetrable taste in music. There have also been many others who have made CMU home for these last four years and for that I am truly grateful to: Debbie Basu, Dan Clymer, Ross Cunningham, Prathamesh Desai, Jason Fox, Zack

vi

Francis, Stephen Herd, Nick Jones, Christopher Kantzos, Varun Kasireddy, Edgar Mendoza, Greg Le Mon, Sudhanshu Nahata, Recep Önler, Rahi Patel, Georges Saliba, Oleg Sapunkov, Avinash Siravuru, Jonelle Yu, and Jackson Wirekoh.

Finally, I am deeply indebted to the many teachers and mentors who have guided me over the years: Mrs. Kristine Alvarez, Mr. John Cotter, Prof. Oscar Crisalle, Mrs. Mary Dumas, Ricky Forquer, Prof. Karl Gugel, Mr. Eric Kincaid, Mrs. Jane Larke, Mark Lusk, Mrs. Kim Mellie, Sra. Michelle Leversee, Don Ohi, Steve Raque, Diane Raque, Prof. Sergei Shabanov, Ralph Utzman, and Brian Yoo. Most of all, I owe my curiosity, my ethics, my drive, and my happiness to my parents Earl and Joy Lynn Scime, my brother Ethan Scime, my grandparents Earl and Leigh Ann Vogel and John and Delite Piper, and my dear friend Phil Tucker.

Abstract

Metal Additive Manufacturing (AM) promises an era of highly flexible part production, replete with unprecedented levels of design freedom and inherently short supply chains. But as AM transitions from a technology primarily used for prototyping to a viable manufacturing method, many challenges must first be met before these dreams can become reality. In order for machine users to continue pushing the design envelope, process space must be expanded beyond the limits currently recommended by the machine manufacturers. Furthermore, as usable process space expands and demands for reduced operator burden and mission-critical parts increase, in-situ monitoring of the processes will become a greater necessity.

Processing space includes both the parameters (e.g. laser beam power and travel velocity) and the feedstock used to build a part. The correlation between process parameters and process outcomes such as melt pool geometry, melt pool variability, and defects should be understood by machine users to allow for increased design freedom and ensure part quality. In this work, an investigation of the AISi10Mg alloy in a Laser Powder Bed Fusion (L-PBF) process is used as a case study to address this challenge. Increasing the range (processing space) of available feedstocks beyond those vetted by the machine manufacturers has the potential to reduce costs and reassure industries sensitive to volatile global supply chains. In this work, four non-standard metal powders are successfully used to build parts in an L-PBF process. The build quality is compared to that of a standard powder (supplied by the machine manufacturer), and correlations are found between the mean powder particle diameters and as-built part quality.

As user-custom parameters and feedstocks proliferate, an increased degree of process outcome variability can be expected, further increasing the need for non-destructive quality

viii

assurance and the implementation of closed-loop control schema. This work presents two Machine Learning-based Computer Vision algorithms capable of autonomously detecting and classifying anomalies during the powder spreading stage of L-PBF processes. While initially developed to serve as the monitoring component in a feedback control system, the final algorithm is also a powerful data analytics tool – enabling the study of build failures and the effects of fusion processing parameters on powder spreading. Importantly, many troubling defects (such as porosity) in AM parts are too small to be detected by monitoring the entire powder bed; for this reason, an autonomous method for detecting changes in melt pool morphology via a high speed camera is presented. Finally, Machine Learning techniques are applied to the in-situ melt pool morphology data to enable the study of melt pool behavior during fusion of non-bulk part geometries.

Table of Contents

Dedicationiii
Acknowledgmentsiv
Abstractviii
Table of Contentsx
Table of Figuresxvii
Table of Tables xxxv
Nomenclaturexxxvii
1 Introduction 1
1.1 Additive Manufacturing (AM)1
1.2 Process Space
1.3 Laser Powder Bed Fusion (L-PBF)7
1.4 Machine Learning (ML) 10
1.5 Motivation 12
1.6 Organization and Major Contributions14
1.7 Formatting Standards19
2 Topic 1: An Exploration of AlSi10Mg L-PBF Process Space
2.1 Background and Literature Review
2.2 Experimental Design and Methods 22
2.2.1 Build Conditions 22
2.2.2 Sample Preparation25
2.2.3 OLP Measurement Techniques
2.2.4 MLP Measurement Techniques

2.3 R	esults	32
2.3.1	Melt Pool Geometry	32
2.3.2	Distribution of Melt Pool Geometries	35
2.3.3	Variability of Melt Pool Geometry across Process Space	41
2.3.4	Bulk Porosity in As-Built MLP Samples	45
2.3.5	Edge Roughness of As-Built MLP Samples	48
2.3.6	A Robust Processing Window	50
2.4 D	iscussion and Summary	51
3 Topic	2: The Effect of Non-Standard Metal Powders on L-PBF Part and Process Quality	55
3.1 Ba	ackground and Literature Review	55
3.2 Ex	xperimental Design and Methods	58
3.2.1	Feedstock	58
3.2.2	Build Conditions	59
3.2.3	Sample Preparation	62
3.2.4	Cylindrical Sample Measurement Techniques	63
3.3 R	esults	67
3.3.1	Qualitative Build Quality and Observations	67
3.3.2	Quantitative Analysis of the Powder Spreading Process	70
3.3.3	Bulk Porosity of As-Built Cylindrical Samples	72
3.3.4	Edge Roughness of As-Built Cylindrical Samples	75
3.3.5	Correlations between Build Quality and Powder Particle Size	77
3.4 D	iscussion and Summary	80

4	Topic 3	3: Autonomous Powder Bed Anomaly Detection and Classification in an L-PB	F
Pro	ocess usin	ng Machine Learning Techniques	2
4	4.1 Bad	ckground and Literature Review8	2
4	4.2 Exp	perimental Design and Methods	7
	4.2.1	Programming Environment8	7
	4.2.2	Powder Bed Camera	7
	4.2.3	Image Pre-Processing	3
	4.2.4	Powder Bed Anomaly Types9	C
	4.2.5	Extraction of Part Geometry Information9	3
2	4.3 Bag	g of Words (BoW) Methodology and Theory9	6
	4.3.1	Overview9	6
	4.3.2	Selection of the Training Data9	7
	4.3.3	Filter Bank	3
	4.3.4	Training	C
	4.3.5	Patch-wise Classification	1
	4.3.6	Layer-wise Classification and Contextual Heuristics	3
2	4.4 Mu	ulti-scale Convolutional Neural Network (MsCNN) Methodology and Theory 10	5
	4.4.1	Overview	5
	4.4.2	Selection of the Training Data10	6
	4.4.3	Input Layer	3
	4.4.4	Hidden Layers11	C
	4.4.5	Patch-wise Classification	5
	4.4.6	Training	5

4.4.7	Layer-wise Classification	121
4.5 Per	formance of the ML Methodologies	122
4.5.1	Confusion Matrices	122
4.5.2	Validation Accuracy	123
4.5.3	Testing Accuracy	126
4.5.4	Computational Burden	132
4.5.5	Comparison of the Machine Learning Methodologies	134
4.5.6	Portability between Material Systems and L-PBF Machines	135
4.6 Res	sults	137
4.6.1	Case Study Overview	137
4.6.2	Analysis of a Single Layer	138
4.6.3	Global and Local Build Reports	140
4.6.4	Cumulative Anomaly Classifications	142
4.6.5	3D Renderings	143
4.6.6	Pseudo-Probabilistic Anomaly Classifications	144
4.6.7	Additional Analysis Visualization Modalities	146
4.7 Dis	cussion and Summary	146
5 Topic 3	3: L-PBF Build Analyses using Autonomous Powder Bed Anomaly Detect	ion and
Classification	n	151
5.1 Bac	ckground	151
5.2 Res	sults	152
5.2.1	Overhangs and Sudden Delamination	152
5.2.2	High-Aspect Ratio Components	162

	5.2	.3	Critical Orientations of Thin Wall Structures	167
	5.2	.4	Impact of Fusion Processing Parameters on Powder Layer Deposition	170
	5.2	.5	Material Systems with Distinct Appearances	178
	5.2	.6	Detection of L-PBF Machine Malfunctions	182
	5.3	Pro	posed Anomaly Mitigation Strategies for Real-Time Implementations	184
	5.4	Disc	cussion and Summary	186
6	Тор	oic 4:	Ex-Situ Melt Pool Morphology across Inconel 718 L-PBF Process Space	189
	6.1	Bac	kground and Literature Review	189
	6.2	Exp	erimental Design and Methods	191
	6.2	.1	Build Conditions	191
	6.2	.2	Sample Preparation	197
	6.2	.3	1LSB Measurement Techniques	198
	6.3	Res	ults	201
	6.3	.1	Melt Pool Geometry	201
	6.3	.2	Distribution of Melt Pool Geometries	203
	6.3	.3	Variability of Melt Pool Geometry across Process Space	209
	6.3	.4	A Cursory Analysis of the Effects of a Powder Layer on Melt Pool Geometry	213
	6.3	.5	Melt Pool Morphology Classifications	216
	6.3	.6	A Comparison of In718 and AlSi10Mg Melt Pool Behaviors	219
	6.4	Disc	cussion and Summary	221
7	Тор	oic 4	: Linking In-Situ and Ex-Situ Melt Pool Morphologies using Machine Lear	ning
Te	chniq	ues i	n an L-PBF Process	224
	7.1	Bac	kground and Literature Review	224

7.2	Exp	erimental Design and Methods	231
7.2	2.1	Programming Environment	231
7.2	2.2	High Speed Camera	231
7.2	2.3	Camera Calibration Procedure	233
7.2	2.4	Coaxial Melt Pool Transformation	234
7.2	2.5	Experiment Sequence Overview	237
7.2	2.6	Collection of In-Situ Training Data (Experiment 2)	238
7.2	2.7	Non-Bulk Geometries (Experiments 3a – 3c)	242
7.3	Bag	of Words (BoW) Methodology and Theory	246
7.3	3.1	Overview	246
7.3	3.2	Selection of the Training Data	247
7.3	3.3	Scale Invariant Feature Transform (SIFT)	248
7.3	3.4	Histogram of Oriented Gradients (HOG)	250
7.3	3.5	Multi-Modality Representation of Melt Pool Morphology	252
7.3	3.6	Training	255
7.3	3.7	Melt Pool Classification and Performance of the ML Methodology	260
7.4	Res	ults	261
7.4	4.1	Melt Pool Morphologies across Process Space (Experiment 2)	261
7.4	4.2	Stripe Edges (Experiment 3a)	271
7.4	4.3	Overhang Region (Experiment 3b)	275
7.4	4.4	Contours (Experiment 3c)	277
7.5	Pro	posed Real-Time Implementation Strategies	279
7.6	Disc	cussion and Summary	285

8	Cor	nclusions 2	289
8	3.1	Summary 2	289
8	3.2	Implications 2	293
8	3.3	Future Work	296
Ref	feren	ces3	300
Appendices			316
		Appendix A : Selected AlSi10Mg 0LP Micrographs and Measurements	317
		Appendix B : Selected AlSi10Mg MLP Micrographs	324
		Appendix C : Selected Cylindrical Sample Micrographs	330
		Appendix D : Powder Bed Images Used for Testing	332
		Appendix E : Selected In718 1LSB Micrographs and Measurements	337
		Appendix F : Selected In-Situ High Speed Camera Images	347

Table of Figures

Figure 2.1: An example of a sample from the OLSB experiments
Figure 2.2: An example of a sample from the OLP experiments
Figure 2.3: An example of a sample from the MLP experiments
Figure 2.4: A representative OLP micrograph, specifically from Sample #23 (370 W, 1200 mm/s). Note the half-width and depth measurement notations; the cross-sectional half-area is the region enclosed by the dotted white polygon. The coloration difference between the lower and upper halves of the micrograph demarcates the boundary between the additively manufactured AlSi10Mg substrate and the Al5083 base plate
Figure 2.5: A representative MLP micrograph, specifically from Sample #24 (370 W, 1400 mm/s). The build direction (<i>z</i> -axis) is oriented upwards
Figure 2.6: An example of a selected bulk region, indicated by the gray rectangle overlaid on the binarized version of the micrograph shown in Figure 2.5. The build direction (<i>z</i> -axis) is oriented upwards
Figure 2.7: A representative MLP micrograph, specifically from Sample #20 (370 W, 600 mm/s). The build direction (z-axis) is oriented upwards
Figure 2.8: An example of a selected edge region, indicated by the gray rectangle overlaid or the binarized version of the micrograph shown in Figure 2.7. Note that the internal porosity has been automatically removed. The build direction (<i>z</i> -axis) is oriented upwards
Figure 2.9: Process map of the cross-sectional melt pool width , developed from the OLF experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend
Figure 2.10: Process map of the cross-sectional melt pool depth , developed from the OLF experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend
Figure 2.11: Process map of the cross-sectional melt pool area , developed from the OLF experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend
Figure 2.12: Process map of the cross-sectional melt pool aspect ratio , developed from the OLF experiment data. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend
Figure 2.13: A comparison between the OLSB above-view melt pool widths reported by Narra $\begin{bmatrix} 12 & p & 96 \end{bmatrix}$ (dashed lines) and the cross-sectional (C-S) OLP melt pool widths reported in Figure

Figure 2.15: Normal probability plot of the measured **widths** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.... 38

Figure 2.18: Normal probability plot of the measured **depths** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.... 39

Figure 2.21: Normal probability plot of the measured **area** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.... 40

Figure 3.8: The test artifacts (NIST parts) built with the five Powder Systems listed in Table 3.1.

Figure 3.9: A representative cylinder built with PS #5. The EOS M290 spontaneously failed to raise the powder dispenser (Figure 1.4), resulting in several layers of insufficient powder

Figure 3.10: Cracking (contiguous red oval) and delamination (dashed red oval) during the PS #4 build......69

Figure 4.7: Flowchart of the implementation of the BoW ML technique discussed in this section. 97

Figure 4.8: A graphical representation of 2D convolution. In this example, a vertical edge *filter* is

Figure 4.12: A *confusion matrix* developed from the *validation* dataset and reporting the percentage of *patches* classified correctly by the **MsCNN algorithm** for each anomaly type. . 125

Figure 4.14: A *confusion matrix* developed from the *testing* dataset and reporting the percentage of pixels classified correctly by the **BoW methodology** for each anomaly type. ... 129

Figure 4.15: A *confusion matrix* developed from the *testing* dataset and reporting the percentage of pixels classified correctly by the **CNN methodology** for each anomaly type..... 129

Figure 4.16: A *confusion matrix* developed from the *testing* dataset and reporting the percentage of pixels classified correctly by the **MsCNN methodology** for each anomaly type.129

Figure 4.23: Layer 2,709 (54.18 mm above the build plate) with the powder bed anomalies classified by the MsCNN. The green pixels indicate the CAD outline of the parts at this layer. 140

Figure 5.15: *Local build report* and vertical part profile for a vertically-oriented tensile bar.... 166

Figure 5.22: A schematic representation of the two extreme impeller blade orientations. In the case of the Type A blade orientation, any plastic deformation due to a recoater strike occurs in the same direction as the geometry shift in the subsequent layer. In the case of the Type B blade orientation, the opposite is true – plastic deformation occurs in the opposite direction as the geometry shift in the subsequent layer. As a result, the Type B impeller blades are less likely to recover from a recoater strike, as the following layer will likely be largely deposited on a bed of unfused powder, with no substantial connection to the previous layer of the impeller...... 169

Figure 5.29: Average detections of *super-elevation* and *part damage* throughout the build height as a function of layer-wise energy density. Each data point represents one of the 30 sample blocks. Note that the EOS nominal process parameters for In718 using 20 µm powder layers (and the corresponding layer-wise energy density) are not available to the author..... 175

Figure 5.32: A color image taken of the bronze powder after completion of the first build supported by Matthews International Corporation. Note the characteristic yellow/gold color.

Figure 5.39: Layer 668 (13.36 mm above the build plate). Severe *part damage* is indicated in magenta. At this layer, the powder bed is approximately 1 mm lower than the powder dispenser (versus the nominal gap of 20 μ m). As a result, insufficient powder is fetched from the powder dispenser to completely cover the build area. Indeed, the maximum horizontal (*x*-axis) extent of the powder spread is visible as the vertical line of *debris* classifications in white.

Figure 7.1: Simplified representations of three different non-bulk geometries. (a) A stripe decreasing in width as it intersects the edge of a part at an angle. As the width decreases, residual heating from adjacent melt tracks [184] is expected to influence the morphology of the melt pool. (b) An unsupported overhang. As the melt pool passes over the bed of un-fused powder, the differing thermal conditions [36], [100] are expected to influence the morphology of the melt pool. (c) A contour pass conforming to the exterior geometry of a part. The differing thermal conditions near the edge of the part [100] are expected to affect the morphology of the melt pool. (c) A contour pass conforming to the exterior geometry of a part. The differing thermal conditions near the edge of the part [100] are expected to affect the morphology of the melt pool. (c) A contour pass conforming to the exterior geometry of a part. The differing thermal conditions near the edge of the part [100] are expected to affect the morphology of the melt pool.

Figure 7.2: The high speed camera setup mounted to CMU's EOS M290; Figure 1.4 shows this setup schematically. This image of the high speed camera was taken by Brian Fisher of CMU.232

Figure 7.4: The rectangular fiducial as imaged by the Alicona Infinite-Focus microscope...... 233

Figure 7.6: The melt pool image from Figure 7.5 after transformation. The reference frame is coaxial, i.e. Lagrangian. The nose, tail, and centerline (dotted line) of the melt pool are Figure 7.7: Eight subsequent high speed camera frames, transformed such that the melt pool appears to be spatially stationary and varying temporally. The relative size scale is preserved Figure 7.8: A CAD rendering of the E2 rectangular exposure region built up over many sets of Figure 7.10: A CAD rendering of the completed E3a test artifact. Note the FoV of the high speed Figure 7.12: A CAD rendering of the completed E3b test artifact. Note the FoV of the high speed Figure 7.14: A CAD rendering of the completed E3c test artifact. Where Δ is the distance from the contour melt track to the edge of the part. Note the FoV of the high speed camera....... 246 Figure 7.16: Flowchart of the implementation of the BoW ML technique discussed in this section. The representative micrographs shown in step (i) are also shown in Figure 6.27. The human "stick figure" shown in between steps (h) and (i) was inspired by the XKCD web comic series created by Randall Munroe [211]......247 Figure 7.17: An example false-color and coaxially-transformed image of a melt pool produced Figure 7.18: The SIFT *features* extracted from the melt pool shown in Figure 7.17. This visualization of the SIFT features was generated using a MATLAB script written by Dollar [214]. Figure 7.19: False-color and coaxially-transformed images of a melt pool produced using PV #1. The top row shows the melt pool image after three different contrast adjustments. The middle row shows each contrast-adjusted image segmented into spatter, the nose region, and the tail

Figure 7.23: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *desirable* ex-situ outcomes.263

Figure 7.33: The melt pool morphologies produced by four different process parmeter combinations during E2. PV #1 is the EOS nominal parameter combination. PV #36 produced the largest percentage of *balling* melt pools. PV #4 produced the largest percentage of *undermelting* melt pools. PV #19 produced the largest percentage of *keyholing porosity* melt pools. 271

Table of Tables

Table 2.1: Process parameter combinations used for each OLP and MLP sample on the EOSM280 L-PBF machine as well as the EOS-nominal parameter combination
Table 2.2: Corresponding measurement errors for the OLP experiments
Table 2.3: Comparison between power and linear fits of AlSi10Mg process map data
Table 2.4: Goodness-of-fit between the measured melt pool distributions and their equivalent normal distributions. R^2 values less than 0.80 are highlighted for convenience
Table 3.1: Nominal Powder System characteristics as reported by the powder manufacturers. 58
Table 3.2: Volume-weighted Powder System characteristics as measured by collaborators atCMU and NCSU
Table 3.3: Processing parameters used for each build on the EOS M290 L-PBF machine
Table 3.4: The furnace profiles used for the stress-relief heat treatments
Table 3.5: Results from an application of the Pearson rank correlation test to the quantitative analysis of the powder spreading process. Table cells are shaded progressively darker for p-value cutoffs of 0.20, 0.10, and 0.05
Table 3.6: Results from an application of the Pearson rank correlation test to the measured bulkporosity types. Table cells are shaded progressively darker for p-value cutoffs of 0.20, 0.10, and0.05
Table 3.7: Results from an application of the Pearson rank correlation test to the measurededge roughness metrics. Table cells are shaded progressively darker for p-value cutoffs of 0.20,0.10, and 0.05
Table 4.1: Descriptions of powder bed anomaly classifications and their respective color codes used throughout this chapter and Chapter 5. Additional descriptions of the anomaly types can be found in Section 4.2.4. Referring to the L-PBF background information provided in Section 1.1 may be useful
Table 4.2: A brief description of the composition of the <i>filter bank</i>
Table 4.3: Overall testing data classification accuracy metrics
Table 4.4: Computer system specifications
Table 4.5: Approximate computation time for selected layer-wise operations
Table 5.1: Process parameter combinations for the AlSi10Mg tensile specimens
Table 5.2: Background powder bed anomaly detections for selected material systems. 181
Table 6.1: Process parameter combinations used for each 1LSB melt track on the EOS M290 L-PBF machine
Table 6.2: Corresponding measurement errors for the 1LSB experiments

Table 6.3: Goodness-of-fit metrics for the In718 melt pool geometry process maps
Table 6.4: Goodness-of-fit between the measured melt pool distributions and their equivalent normal distributions. R^2 values less than 0.80 are highlighted
Table 6.5: A quantification of the effect of a 70 μm thick powder layer on melt pool depth for In718
Table 7.1: Process parameter combinations used for each set of contours imaged during E3c.
Nomenclature

Acronyms			
OLSB	Zero Layers of powder, Single Bead		
1LSB	One Layer of powder, Single Bead		
OLP	Zero Layers of powder, Pad		
AlSi10Mg	Aluminum – 10 wt-% Silicon – 0.3 wt-% Magnesium		
AM	Additive Manufacturing		
BoW	Bag of Words		
$C_2H_2O_4$	oxalic acid		
CAD	Computer Aided Design		
CMU	Carnegie Mellon University		
CNN	Convolutional Neural Network		
CONV	Convolution layer		
CPU	Central Processing Unit		
C-S	Cross-Sectional		
CT	Computed Tomography		
CV	Computer Vision		
DED	Direct Energy Deposition		
DEM	Discrete Element Modeling		
DoG	Difference of a Gaussian		
DXR	Dynamic X-Ray		
E1	Experiment 1 (melt pool morphology)		
E2	Experiment 2 (melt pool morphology)		
E3	Experiment 3 (melt pool morphology)		
EBAM	Electron Beam Additive Manufacturing		
EB-PBF	Electron Beam Powder Bed Fusion		
EDM	Electrical Discharge Machining		
FC	Fully Connected layer		
FEM	Finite Element Modeling		
FoV	Field of View		
FPGA	Field Programmable Gate Array		
GD	Gradient Descent		
GPU	Graphics Processing Unit		
H_2O_2	hydrogen peroxide		
HDH	Hydride-dehydride		
HBF ₄	fluoroboric acid		
HOG	Histogram of Oriented Gradients		
LENS	Laser Engineered Net Shaping		
In625	Inconel 625		
In718	Inconel /18		
L-PBF	Laser Powder Bed Fusion		
LRN	Local Response Normalization layer		
ML	Niachine Learning		

- MLP Multiple Layers of powder, Pad
- MsCNN Multi-scale Convolutional Neural Network
 - MSE Mean Square Error
 - NCSU North Carolina State University
 - NIST National Institute of Standards and Technology
 - PCA Principle Component Analysis
 - PH Precipitation Hardened
 - PREP Plasma Rotating Electrode Powder
 - P-S Plasma Spherodization
 - PS Powder System
 - PSD Powder (Particle) Size Distribution
 - PV beam Power and travel Velocity
 - RAM Random-Access Memory
 - ReLU Rectified Linear Unit layer
 - RF Random Forest
 - SGD Stochastic Gradient Descent
- SGDM Stochastic Gradient Descent with Momentum
 - SIFT Scale Invariant Feature Transform
 - SS Stainless Steel
 - SVM Support Vector Machine
- Ti64 Titanium 6 wt-% Aluminum 4 wt-% Vanadium
- t-SNE t-distributed Stochastic Neighbor Embedding
- VLAD Vector of Locally Aggregated Descriptors

Machine Learning and Computer Vision Terminology

confusion matrix	a common performance metric for ML algorithms which compares ix classifications produced by the algorithm to the "ground tru classifications of a human				
dictionary	stores a set of learned visual words				
feature	e a number(s) which describes one or more aspects of the input data				
filter	a function which is convolved with an input data volume in order to extract a <i>feature</i>				
filter bank	k a set of <i>filters</i>				
fingerprint	a vector of numbers which "fully" describes an input data volume				
kernel	another word for a <i>filter</i> , more commonly used in the context of CNN architecture and typically referring to a 2.5D or 3D <i>filter</i>				
morphology	a set of <i>fingerprints</i> describing a set of in-situ melt pool appearance which are similar to each other				
a sub-region of a powder bed image used for training and a patch classification					

response	the value produced by a neuron, e.g. the output of a convolution operation between a <i>filter</i> and an input data volume		
testing data	the set of data which is never used during training and is used to objectively determine the performance of the final ML algorithm		
training data	the subset of the training database to which the ML model is fit		
validation data	the subset of the training database which is used to evaluate the performance of the trained (but not final) ML model and inform the tuning of hyperparameters		
visual word	an average of a set of <i>features</i> ; a set of <i>visual words</i> can be used to effectively describe an input data volume		
weight	the element-wise values of a kernel or a filter		

Powder Bed Anomaly Classification Visualization Modalities

rowaer bea Anomaly elassification visualization modulities			
global build report	reports the anomaly classifications throughout the build height over the entire powder bed		
heat map	reports the cumulative anomaly classifications (over the course of the entire build) across the powder bed		
local build report	reports the anomaly classifications throughout the build height over a sub-region of the powder bed		

Powder Bed Anomaly Classes

okav	no significant anomalies in the powder bed			
debris	debris or other small to mid-sized discrepancies located in the powder bed but not directly over any parts			
disturbance	encompasses both the <i>debris</i> and <i>part damage</i> anomaly classes			
incomplete spreading	occurs when an insufficient amount of powder is repeatedly fetched from the powder dispenser; results in a lack of powder, the severity of which is highest nearest the powder collector			
part damage	general classification for any significant damage to a part; characterized by a variety of signatures			
recoater hopping	caused by the recoater blade striking a part, characterized by repeated vertical (parallel to the <i>y</i> -axis) lines			
recoater streaking	caused either by the recoater blade dragging a contaminant across the powder bed or by damage to the blade; characterized by horizontal (parallel to the <i>x</i> -axis) lines			
super-elevation	occurs when a part warps or curls upwards out of the powder layer; typically the result of a buildup of residual thermal stresses or swelling			

Melt Pool Morphology Anomaly Classes

ment i con merphology / moniary classes			
balling	a surface tension-driven anomaly which causes a melt pool to break up into discrete "balls" of material		
desirable	a melt pool which does not produce any observable ex-situ defects		
keyholing porosity	a melt pool which produces porosity due to the unstable collapse of the keyhole vapor cavity		
severe keyholing	defined (in this thesis) as a melt pool with a depth to half-width ratio greater than 2.5; note that no distinct in-situ appearance is observed for this type of melt pool		
spatter	a melt pool with a <i>fingerprint</i> (in-situ appearance) dominated by the presence of hot (and therefore visible) ejecta from the melt pool		
under-melting	a melt pool with a depth less than that of the effective layer thickness; note that this does not necessarily produce lack-of-fusion porosity		

Variables			
A	melt pool cross-sectional area		
α	thermal diffusivity		
χ^2	chi-square statistical test		
c_p	specific heat capacity		
Δ	the distance of a melt track from the nominal edge of a part during a		
$E(\Omega)$	loss function		
$\int f$	frame rate of the high speed camera		
F	size of a <i>filter</i>		
Fo	Fourier Number		
γ	momentum coefficient		
Г	image contrast adjustment		
h	hatch spacing		
Н	the distance from the x-y plane down to the top of the build		
Ι	pixel value in a high speed camera image		
η	learning rate		
k	thermal conductivity		
Х	powder consolidation factor		
l_e	effective powder layer thickness		
l_n	nominal powder layer thickness		
L_c	critical length during thermal diffusion		
L _{fov}	length of an E2 melt track within the FoV of the high speed camera		
L_{tot}	total length of an E2 melt track		
Ν	number of E2 melt tracks in a set		
ω	kernel weight		
Ω	set of <i>kernel weights</i> throughout the depth of the CNN		

laser beam power Ρ

Р extents of the padding during a convolution operation

φ response of a kernel (neuron)

Q absorbed laser beam power

 R^2 **R-square statistical metric**

ρ correlation coefficient

ρ density

layer-wise energy density и

r distance from the beam spot in a Lagrangian reference frame

edge roughness metric: arithmetic average of absolute values Ra

edge roughness metric: root mean squared Rq

Rsk edge roughness metric: skewness

Rz edge roughness metric: maximum peak-to-valley difference

S stride of a *filter*

ξ distance parallel to the laser beam travel direction

a measure of time during thermal diffusion t

time over which high speed camera data are collected during E2 t_{data}

time during E2 that the melt pool is within the FoV of the high speed t_{inframe} camera

t_{maxrecord} t_{tot}

 $T(\xi, x, y)$

maximum time the high speed camera can record data total time required to expose a set of 1LSB melt tracks during E2

melt pool isotherm temperature in a Lagrangian reference frame

- local background temperature T_0
- laser beam travel velocity v
- W_i input size of a data volume in a CNN

output size of a data volume in a CNN W_{o}

width of a set of E2 1LSB melt tracks W_{tot}

distance anti-parallel to the direction of the recoater blade х

distance perpendicular to x and zy

z distance in the vertical (build) direction

1 Introduction

1.1 Additive Manufacturing (AM)

As its name suggests, Additive Manufacturing (AM), colloquially known as 3D Printing, allows for the production of parts by successively adding layers of raw material on top of each other. This is in direct contrast to more traditional, subtractive, forms of manufacturing; a contrast that results in a wealth of new capabilities as well as many challenges which must first be addressed before these capabilities can be fully realized. Additive Manufacturing spans an incredibly wide range of materials and processes (Figure 1.1) and has been under development since 1982 [1]. Over the past two decades, AM technologies capable of producing metal parts have been transitioning from prototyping tools to viable manufacturing methods [2].

This transition has been driven by increasingly favorable economics [3], [4] as well as the desire of companies, particularly in the aerospace, biomedical, energy, automotive, and tooling sectors [5]–[7], to exploit the increased design freedom that this technology promises [3], [8]. Design freedom in this context refers to not only the physical design of a part, but also the processing parameters that are used to build the part, and even the type of feedstock used as raw material. This freedom can allow for the creation of complex internal features such as cooling channels and lattice structures [2], [3], [9], production of assemblies as single parts [10], and monolithic parts with engineered material property gradients [11], [12]. Additionally, the opportunity to dramatically shorten the supply chain is inherent to AM [13]–[15] which could make additive technology ideal for producing mission-critical parts, on-demand, at remote locations [2], [16].





Figure 1.1: The breadth of Additive Manufacturing technologies, including methods such as vat photopolymerization, material extrusion, material jetting, binder jetting, powder bed fusion, direct energy deposition, and sheet lamination [17]⁶. The blow-up region of this figure focuses on the metal AM processes.

⁶ Prior approval for use of this figure was obtained by the author from 3D Hubs' copyright department.

Most metal AM processes can be broadly grouped into three categories: Binder Jetting (e.g. ExOne[™] and Desktop Metal[™]), Direct Energy Deposition (DED), and Powder Bed Fusion (PBF). DED can be subdivided into Laser Engineering Net Shape (LENS) powder stream processes (e.g. OPTOMEC[®]), Electron Beam Additive Manufacturing (EBAM) wire feed processes (e.g. Sciaky Inc.), laser hot wire processes (e.g. Lincoln Electric[®]), and robotic arcwelding process (e.g. TWI Ltd). Similarly, PBF can be subdivided into technologies that rely on a laser to melt the powder (e.g. EOS GmbH, SLM Solutions GmbH, ConceptLaser GmbH, and Renishaw[®]) and technologies that use an electron beam to perform the melting (e.g. Arcam[®]). It is worth noting that as metal AM continues to gain market acceptance, new processes are being developed which fall outside of these core technologies, such as Aerosol Jet Printing (e.g. OPTOMEC[®]) for electronics applications and magnetically-controlled molten metal droplet deposition (e.g. Vader Systems) for aluminum components.

In Binder Jetting, unlike the other metal AM processes mentioned above, no melting occurs during the actual layer-wise printing operations; instead, sintering or infiltration of the metal powders occurs as a post-processing step [18], [19]. Some Binder Jetting machines operate by depositing sub-millimeter sized droplets of binder (i.e. glue) onto a bed of metal powder [18] while others extrude slurry of metal particles entrained within the binder [19]. After printing is complete, the part is considered "green" and one or more heat treatment cycles are used to harden the binder and then either sinter the metal powder particles or infiltrate the green part with a low melting temperature alloy such as bronze [18].

3

The LENS process is derived from laser cladding technology and uses a laser beam to form a melt pool on a substrate; as the laser moves, material is added to the melt pool using a stream of metal powder entrained within a jet of inert gas [20]. The EBAM process uses an electron beam to form a melt pool into which material is added in the form of one or more metal wires [21]. Laser hot wire processes are relative newcomers to the AM arena and use a laser beam to form a melt pool into which a resistively pre-heated metal wire is fed [22]. DED AM also encompasses more traditional welding processes, such as arc welding, which are becoming increasingly automated [23]. DED processes offer a number of advantages, including high build rates [5], large build volumes [21], and in some cases the ability to build with multiple alloys simultaneously [11], [21]. However, parts produced using DED technologies are generally only considered to be "near net-shape" and often require extensive post-processing before use [5].

PBF processes operate by using either an electron beam or a laser beam to selectively melt regions of a bed of metal powder [24], [25]. Electron Beam PBF (EB-PBF) systems typically operate at higher build chamber temperatures and have higher material deposition rates than Laser PBF (L-PBF) machines [26]. Conversely, L-PBF generally offers higher achievable feature resolution and a wider range of available material systems than EB-PBF [27], [28]. L-PBF is described in further detail later in this chapter.

1.2 Process Space

The material deposition behavior of DED and PBF technologies can be compared through the visualization of a "process space" [29] defined by the power and the travel velocity (PV) of their heat sources (i.e. the "beam"), as shown in Figure 1.2. It is readily apparent from Figure 1.2 that the electron beam processes generally operate at higher powers than their laser counterparts while the PBF processes generally operate at higher beam velocities than the DED processes.



Figure 1.2: PBF and DED processes plotted in process space. This figure was originally developed by Beuth et al. [29] and has been updated by Fox [30, Fig. 3], Francis [31, Fig. 1.1], Gregory Le Mon of Carnegie Mellon University, and the author to more accurately reflect the current state of the industry.

The concept of process space, as defined by the beam power and travel velocity, is used extensively throughout this thesis to motivate the understanding of melt pool geometry behavior and delineate process defect regimes. Specifically, "process mapping" is a technique developed by Beuth et al. [32] which defines curves of constant process outcomes (e.g. melt pool geometry and as-built microstructure) with respect to critical process parameters (e.g. beam power and travel velocity). These curves can then be used to describe the behavior of other process outcomes such as porosity, mechanical properties, and required build time [33]. In the stylized L-PBF process map shown in Figure 1.3, all of the beam power and velocity combinations which fall on the solid red line are expected to produce melt pools with similar cross-sectional areas (see the following section) while all of the beam power and velocity combinations which fall within the highlighted red region are expected to produce melt pools which experience keyhole-mode melting conditions.



Beam Velocity

Figure 1.3: A stylization of L-PBF process space overlaid by three curves of constant melt pool cross-sectional area where $A_1 > A_2 > A_3$ and one curve of constant melt pool length to cross-sectional width ratio. Some of the curves of constant melt pool geometry delineate regimes of process space for which the formation of certain processing defects can be expected. More details regarding these processing defects and desirable processing windows can be found in Chapters 2 and 6.

While dozens of process parameters impact the quality and performance of L-PBF manufactured parts, six are considered to have a controlling influence on melt pool geometry and as-built microstructure [29], [31], [32]. A common approach is to first study the influence of (1) beam power and (2) beam travel velocity for a given DED or PBF process and material system before performing targeted studies of the remaining critical process parameters. For example, work by Fisher [34], [35] demonstrates that melt pool size and the size scale of microstructural features increase as (3) background temperature⁷ increases. Francis [31] found that the (4) beam diameter can dramatically affect the shape of the melt pool while Montgomery [36] found that melt pool size increases with increasing (5) powder layer thicknesses. Finally, Fox [37] and Chapter 7 of this thesis study the influence of (6) local feature geometry on melt pool behavior and flaw formation. This work also considers the feedstocks (e.g. powders and wires) used in AM processes to be part of overall process space – a concept which is explored throughout Chapter 3.

1.3 Laser Powder Bed Fusion (L-PBF)

The work herein focuses on the L-PBF process – currently one of the mostly widely deployed and industry-relevant AM technologies [38]. This process operates by spreading a thin layer, typically 20 μ m to 120 μ m thick, of metal powder over a build plate using a recoater blade. After powder spreading, a laser beam is used to selectively melt the powder in locations corresponding to a 2D slice of a 3D part. The locally-molten region of the powder bed is

⁷ In this context, the term "background temperature" refers to the temperature of the material surrounding the melt pool. While this temperature is often directly related to the temperature of the build chamber, it can also be influenced by previous melt tracks [184] and previous layers [224].

typically referred to as a "melt pool." After the lasing is complete, the build plate is lowered, another layer of powder is spread (now over an existing powder bed), and the process repeats until the part is finished. The entire process of creating a part is often referred to as a "build" and occurs within a build chamber purged with an inert gas such as argon or nitrogen at a typical absolute pressure of approximately 1 atm [39]. A schematic representation of the EOS M290 L-PBF machine (EOS GmbH, Germany) [40] at Carnegie Mellon University's (CMU) NextManufacturing Center is shown in Figure 1.4; this figure includes components and modifications relevant to this work and which will be discussed in detail, as appropriate, throughout this document. Figure 1.5 shows an image of the powder bed, taken from inside the EOS M290's build chamber; the EOS M290 build plate measures 250 mm × 250 mm and the build volume extends to 325 mm in height [40]. The EOS M290 machine has a maximum nominal beam power of 370 W and a nominal D86 beam diameter of 100 µm [40].



Figure 1.4: A schematic representation of the EOS M290 machine at CMU's NextManufacturing Center. The arrows indicate the direction that the schematic components will move immediately following the lasing of the current layer.



Figure 1.5: An image taken from inside of the EOS M290 build chamber using the "powder bed camera" shown in Figure 1.4. It is oriented such that it is looking down (negative *z*, Figure 1.4) at the powder bed.

Figure 1.6 shows a 3D Finite Element Model (FEM) of a melt pool for purely illustrative purposes. A schematic of two adjacent melt tracks is shown in Figure 1.7. L-PBF machines fuse each powder layer according to a prescribed "scan strategy" that governs the laser beam travel path. The default EOS M290 scan strategy [41] is shown schematically in Figure 1.8 and fuses the powder bed in a sequence of "stripes." Each stripe is formed by rastering the laser beam across the width of the stripe, with each adjacent melt track separated by a distance referred to as the hatch spacing. For the EOS M290 it is also standard to apply one or more "contour passes" to the perimeters (both internal and external) of each part. During a contour pass, the laser beam path is conformal to the perimeter geometry of the part and is offset inwards by a distance that typically ranges between 0 µm and 100 µm. Additional scan strategies are implemented by other L-PBF machines but are not discussed in this thesis.



vidth overlap depth area

Figure 1.6: A reference image taken from a 3D FEM to illustrate a melt pool. Specifically, this is an ABAQUS model based on the work of Solyemez [42] and is of an L-PBF-processed AlSi10Mg melt pool at an absorbed power of 200 W and a beam travel velocity of 1200 mm/s. The molten region is delineated by the colored (non-gray) elements which are a temperature greater than the liquidus temperature. Note that the melt pool is sectioned along its ξ -axis line of symmetry. Unlike the *x*-axis, the ξ -axis is always parallel to the travel direction of the laser beam with an origin centered at the beam spot.

Figure 1.7: A cross-sectional view of two idealized adjacent melt tracks with key dimensions annotated. These dimensions are referred to extensively in Chapters 2 and 6. This figure is based on a similar figure presented by Narra [12, Fig. 5.8].



Figure 1.8: A schematic of the default EOS M290 scan strategy. The powder layer is fused with a sequence of stripes, each of which is composed of adjacent melt tracks separated by the hatch spacing. The orientation of the scan strategy (i.e. the stripes) rotates by 67° every layer in order to improve inter-layer bonding [41].

1.4 Machine Learning (ML)

The L-PBF process operates over an immense range of size and time scales. For example, while the melt pool and many defects are on the order of tens to hundreds of microns in size (and form in tens of microseconds), the laser beam path may meander for tens of kilometers (over a period of a week of more) within a single part. Effectively monitoring such behavior requires the analysis of complex and often poorly-understood datasets. Furthermore, as AM technologies evolve rapidly, so must any monitoring techniques. In other words, requiring a human programmer to substantially redesign a process monitoring algorithm every time a new AM machine or material system is developed would be unsustainable. Fortunately, the field of Machine Learning (ML) offers a powerful toolset which is well-equipped for tackling these data analysis challenges.

Fundamentally, all ML algorithms operate by extracting *features* from a set of training data provided by a human [43]. The extracted *features* are then analyzed: their frequencies in the training data are quantified and their similarities and differences are described [44]. In the case of Deep Learning, the algorithm will actually design its own optimized set of *feature* extraction tools, as opposed to only using the tools provided by the human programmer [43]. Once the *feature* extraction system is robust, a model for describing the input data (based on its *features*) is created [43]. *Features* are extracted from any new data (i.e. data a user wishes to analyze) and are input into the model allowing the algorithm to make a decision (e.g. whether or not a process anomaly is present) that is informed by the knowledge contained within the training database. Figure 1.9 shows a high-level schematic of a generic ML algorithm. While ML can theoretically be applied to any arbitrary data set, all of the ML algorithms presented in this work operate on visible-light image data, the *features* of which are extracted using well-established Computer Vision (CV) techniques.



Figure 1.9: A schematic of a generic Machine Learning algorithm. Features are extracted from the input data, analyzed, and used to create a model describing the training data. In Deep Learning, the feature extraction step may be optimized based on the efficacy of the model. Once a robust model is complete, the algorithm can make decisions about new, previously unseen, input data.

Finally, it should be noted that the aforementioned span of size and time scales necessitates the collection of highly resolved data over a long period of time. This requirement results in significant data storage and data transmission burdens, particularly if real-time process monitoring is desired. This additional challenge is addressed, where appropriate, throughout this thesis.

1.5 Motivation

The industries and applications for which Additive Manufacturing is most applicable are also industries and applications for which design freedom is highly valued, confidence in a robust supply chain is required, and quality assurance is paramount (Section 1.1). Currently, AM machine manufacturers offer only a limited set of "approved" processing parameters for each material [45], [46], limiting an end-user's ability to design components to, for example, be resistant to fatigue failure [47], [48] or have specified microstructures [12], [49], [50]. PBF machine manufacturers also strongly recommend that their customers use only powders from vetted vendors, produced via specific processes, and with restrictive size distributions [51], which can increase the cost and risk of using the technology [3], [4], [52]. Finally, the required levels of quality assurance and process reliability are difficult to achieve with the systems currently on the market [2] and can likely only be achieved through the implementation of insitu process monitoring and closed-loop control schema [2].

As designers push the boundaries of the processing parameters, it becomes necessary for them to understand how those parameters impact the final part quality and what kinds of defects and degree of variability they can expect. **Topic 1 (Chapter 2)** addresses this challenge

12

with a case study of the effects of processing parameters on process outcomes in the AlSi10Mg aluminum alloy. Now consider a major aerospace corporation looking to protect its supply chain by sourcing powder from multiple vendors, or a United States Navy aircraft carrier docked at a foreign port and looking to take on additional feedstock so that it can continue producing critical replacement parts for aircraft while underway. Both of those events are only possible if the processing parameters can be adjusted to accommodate alternative powders and if their effects on part quality are understood. **Topic 2 (Chapter 3)** investigates this "powder space" through the successful use of four non-standard Powder Systems.

Consider again the situation of additively producing replacement parts for aircraft at sea; it is of paramount importance that those parts be defect-free. Additionally, while a naval vessel may have the manpower to "babysit" an AM machine during printing, a scientific outpost in Antarctica or on Mars almost certainly would not. Real-time, autonomous monitoring of the powder bed has the potential to allow for increased quality assurance, increased process stability, and a reduced operator burden; **Topic 3 (Chapters 4 and 5)** presents a Machine Learning-based Computer Vision algorithm that tackles this challenge. Finally, many of the defects introduced by Topic 1 are too small to detect using the techniques observed by Topic 3 and may still occur while building complex geometries even if the relationships between the processing parameters and the process outcomes are well understood. Because quality assurance is crucial in so many AM applications, **Topic 4 (Chapters 6 and 7)** presents a Machine Learning-based methodology for linking melt pool morphology, captured by a visible-light high speed camera, to process defects. The development of such a methodology is a crucial step on the path toward creating a practical feedback control system. It is the author's hope that the work presented in this thesis helps increase the acceptance of Additive Manufacturing as a viable manufacturing method, with a particular focus on enabling it to produce mission-critical components in remote locations.

1.6 Organization and Major Contributions

This thesis is organized around two key research themes: (1) expanding process space and (2) developing in-situ process monitoring methodologies. Each research theme guides two research topics and each topic is covered by either one or two chapters. Each topical chapter is designed to be self-contained, consisting of relevant background information, a literature review, applicable methods and theory, and a description and discussion of the completed work. A brief overview of each chapter (including the Introduction, Conclusions, and six topical chapters) as well as a preview of some contributions of this work are provided below:

- The first chapter introduces the reader to Additive Manufacturing and its current status in the greater manufacturing world. Motivation for the work presented in this thesis is provided and the structure and format of the document are summarized.
- 2. The second chapter (Topic 1) explores L-PBF processing space for the AlSi10Mg aluminum alloy. Process maps are presented for critical melt pool geometric dimensions. The statistical distribution and variability of melt pool dimensions are analyzed across process space. Bulk porosity (produced by two different mechanisms) and 2D edge roughness are measured across process space. Finally, a robustly-defined process parameter window is proposed. The results presented in this chapter provide

14

designers and machine operators with information critical for choosing appropriate processing parameters for the AlSi10Mg material system.

- 3. The third chapter (Topic 2) expands L-PBF process space by using four non-standard metal Powder Systems to successfully manufacture a standardized set of test artifacts. Note that the largest powders tested included particle sizes up to 2.4 times larger than those found in the manufacturer-recommended powder. For each Powder System, the quality of the powder layers is analyzed and the as-built part quality is quantified by measuring the 2D edge roughness and bulk porosity of the test artifacts. The powder layer and part quality results are compared to those of test artifacts built using a standard Powder System. Finally, while overall quality remained high across the tested Powder Systems, correlations between powder particle size and both powder layer quality and as-built part quality are identified. A portion of the work presented in this chapter is published in an America Makes report titled "A Database Relating Powder Properties to Process Outcomes for Direct Metal AM" [53].
- 4. The fourth chapter (Topic 3) discusses the development of two algorithms designed to autonomously analyze the powder spreading portion of the L-PBF process. Both algorithms apply Machine Learning and Computer Vision techniques to successfully detect and classify several types of millimeter-scale anomalies present on the powder bed. The first algorithm utilizes a "Bag of Words" approach while the second algorithm leverages transfer-learning to train a Convolutional Neural Network to analyze a multi-scale dataset. Critically, the second algorithm methodology does not rely on human-created heuristics and is therefore highly extensible to alternate material systems and

other powder bed-based AM technologies. The performances of the two algorithms are quantified and compared. Finally a case study is used to demonstrate the broad capabilities of the final algorithm. A significant portion of the work presented in this chapter is published in the Additive Manufacturing Journal under the title "Anomaly Detection and Classification in a Laser Powder Bed Additive Manufacturing Process using a Trained Computer Vision Algorithm" [54]. It is also worth noting that the presented Convolutional Neural Network analyzes data at multiple size scales – a technique not found commonly in the existing literature.

- 5. The fifth chapter continues **Topic 3** by applying the final powder bed anomaly detection algorithm in a variety of case studies. In particular, the algorithm is shown to provide insights on the delamination of parts from the build plate, the printing of high aspect ratio, thin wall, and overhanging structures, and the influence of non-standard process parameters on the surrounding powder bed. Finally, the algorithm is shown to perform robustly for non-standard material systems and its hypothetical usage in a real-time environment is briefly discussed. A significant portion of the work presented in this chapter is published in the Additive Manufacturing Journal under the title "Anomaly Detection and Classification in a Laser Powder Bed Additive Manufacturing Process using a Trained Computer Vision Algorithm" [54].
- 6. The sixth chapter (**Topic 4**) develops a database of melt pool geometry in the presence of a powder layer for the Inconel 718 alloy across L-PBF process space. This database is based on ex-situ measurements and is used in the seventh chapter to correlate in-situ data with process outcomes. The size and composition of the database also allowed for

16

the analysis of the statistical distribution and variability of melt pool dimensions across process space. Based on comparisons to data in the literature, a brief discussion of the effects of a powder layer on melt pool dimensions is provided. Finally, melt pool behavior is compared between the Inconel 718 and AlSi10Mg (second chapter) material systems. Significant outliers in measured melt pool cross-sectional size are observed for both material systems and strongly motivate the need for additional work in this area.

7. The seventh chapter is the heart of **Topic 4**. In-situ melt pool morphologies are extracted from high speed camera images using Computer Vision techniques while a "Bag of Words" Machine Learning approach is used to cluster the morphologies. The clusters are then linked to process flaws such as porosity and surface tension instabilities using the ex-situ database constructed in chapter six. This linkage is successful, demonstrating that certain in-situ melt pool morphologies arise only in particular regions of process space. Finally, the CV/ML methodology is applied to the preliminary study of melt pool morphology during the exposure of several non-bulk geometries including the edges of stripes, overhangs, and contours. The ability to study fusion of such geometries could allow for the design of optimized process parameters in the future. Note that much of the work presented in this chapter relies upon the use of a custom algorithm which transforms the collected melt pool data into an Eulerian reference frame – an operation with broad implications for the further study of melt pool dynamics. Some of the work presented in this chapter is published in the Society of Manufacturing Engineers Letters under the title "Using Coordinate Transforms to

17

Improve the Utility of a Fixed Field of View High Speed Camera for Additive Manufacturing Applications" [55].

8. The eighth and final chapter summarizes the results from the six topical chapters and places them within a contextual framework to demonstrate their interdependence and make cross-topic conclusions. The implications of the presented work are discussed in the context of the motivations provided in the first chapter. Finally, significant and exciting future work is identified and motivated.

1.7 Formatting Standards

The Additive Manufacturing terminology used in this document complies as closely as possible with ISO/ASTM 52900:2015 [56]. All of the figures presenting data relating to processing defects are colored either as shown in Figure 1.3 or presented in grayscale. Specifically, keyholing porosity is shown in red, lack-of-fusion porosity is shown in light blue (teal), balling is shown in purple (magenta) and desirable regions of process space are shown in green. The color scheme used to represent melt pool features in Figures 1.6 - 1.8 is also used throughout this document. A standardized, right-handed and self-consistent coordinate system (shown in Figures 1.4 - 1.8) is used throughout this document. Note that some of the images of the powder bed presented in Chapters 4 and 5 have the negative signs omitted from the y-axis for clarity. All of the algorithm architecture schematics for the presented ML methodologies adhere as closely as possible to the color scheme used in Figure 1.9, where input data are shown in gray, feature extraction tasks are shown in purple (magenta), feature analysis tasks are shown in orange, and components of the ML model itself are shown in light blue (teal). All acronyms, italicized words, and variables used in this thesis are defined in the Nomenclature section.

2 Topic 1: An Exploration of AlSi10Mg L-PBF Process Space

2.1 Background and Literature Review

In their endeavors to exploit the design freedoms that come with Additive Manufacturing, designers must push the boundaries of the machine processing parameters. It therefore becomes necessary for them to understand how those parameters impact the final part quality and what kinds of defects and degree of melt pool dimensional variability they can expect. Currently, AM machine manufacturers offer only a limited set of "approved" processing parameters for each material system [45], [46], limiting the design freedom of the end user. Process mapping is a technique developed by Beuth et al. [32] that enables the correlation of process parameters to process outcomes (e.g. melt pool geometry, porosity, and as-built microstructure) by plotting those outcomes across process space.

This work applies the process mapping approach to the EOS AlSi10Mg (Al 87 – 90 wt-%, Si 9.0 - 11.0 wt-%, Mg 0.2 – 0.45 wt-%, Fe \leq 0.55 wt%, Mn \leq 0.45 wt% [57]) aluminum alloy in an L-PBF process. Aluminum alloys are highly attractive to many of the industries implementing Additive Manufacturing as a result of their high thermal conductivity (173 W/m-K at 20 °C for AlSi10Mg [57]) and high strength-to-mass ratio [58]. Relevant applications include lightweight heat exchangers for aircraft and spacecraft [59], [60] and engineered meshes for zero-gravity fuel tanks [61]. Aluminum alloys are notoriously difficult to manufacture additively due to their low absorptivity⁸, high thermal conductivity [62], propensity for vaporization of the aluminum and associated alloying elements [63], and susceptibility to hot-cracking [62]. Similar to

⁸ Approximately 12% at 300 °C and 40% at 1000 °C for a wavelength of 1060 nm (near the output of the EOS M290's Yb: YAG laser [75]) [225, pp. 70–72].

traditional casting alloys, AlSi10Mg is near the Al-Si eutectic composition (i.e. has a narrow melting range of approximately 40 K [64]) which reduces its susceptibility to hot-cracking during rapid solidification [62], [65].

While extensive bodies of process mapping work exist for other alloys [12], [29], [31], [32], [66]–[68], at the time of this research, AlSi10Mg was a relative newcomer to the AM community. Aboulkhair et al. [69] indirectly investigated the effect of beam travel speed on the microstructure of single melt tracks via micro-hardness testing. Read et al. [70] compared creep behavior in additively manufactured AlSi10Mg to traditionally-formed AlSi10Mg. Tang et al. [71] performed extensive work on the fatigue performance of additively manufactured AlSi10Mg and studied lack-of-fusion porosity at several process parameter combinations [72]. Most notably, Kempen et al. [58] worked to find optimal processing parameters for AlSi10Mg in an L-PBF process but within a narrower power range than investigated in the work presented in this chapter. Many of the experiments that this work is based on were also used by Dr. Sneha Prabha Narra of CMU to control the as-built microstructure (cell spacing) of AlSi10Mg components produced via L-PBF [12, Ch. 5].

In this chapter, process maps are presented that correlate beam power and beam travel velocity to cross-sectional melt pool width, depth, area, and aspect ratio. Melt pool width results are compared to those reported by Narra [12, p. 96]. The presented melt pool geometry results are supported by a statistical analysis which also allows for statements to be made concerning the statistical distribution of melt pool dimensions and the variability of melt pool geometry across process space. While Francis [31, Fig. 4.5] and others [73] have characterized the variable behavior of melt pool depth during keyhole-mode melting, the existing literature is

21

relatively sparse with regards to the statistical study of general melt pool variability in the AM processes.

The occurrences of two types of processing defects: lack-of-fusion and keyholing porosity, are investigated across process space. Porosity resulting from the lack-of-fusion mechanism occurs when there is insufficient overlap between melt pools and or the melt pools do not fully penetrate the powder layer (this case is also known as "under-melting") [72], [74]. Keyhole-mode melting occurs in the high energy density (high beam power, low beam velocity) region of process space, where periodic vaporization of the molten material can occur. Under certain conditions, the resultant vapor pocket may become trapped as porosity in the solidified melt pool [74]. Additionally, the as-built surface finish is explored across process space using several 2D edge roughness measures. Finally, a process window based on the quantitative porosity analysis is presented, expanding on the window presented by Narra [12, p. 102]. The work presented in the chapter was supported by the Research for Advance Manufacturing in Pennsylvania program (prime award number FA8650-12-2-7230, sub-award number 543105-78001) and through a donation from Mr. Richard Fieler (Carnegie Institute of Technology class of 1956).

2.2 Experimental Design and Methods

2.2.1 Build Conditions

All of the experiments discussed in this section were performed on a 400 W version of an EOS M280 L-PBF machine at the Arconic (then Alcoa c.a. 2015) Technical Center in New Kensington, Pennsylvania. Three sets of relevant experiments were performed, including single

bead melt tracks with no added powder (OLSB⁹) (Figure 2.1), multiple adjacent melt tracks (19 mm long × 9.5 mm wide) with no added powder (OLP¹⁰) (Figure 2.2), and 19 mm long × 9.5 mm wide × 13 mm tall solid blocks (MLP¹¹) (Figure 2.3). Because it is difficult to procure AlSi10Mg plates an Al5083 base plate was used, with the OLSB and OLP experiments deposited on top of additively manufactured AlSi10Mg substrates. These AlSi10Mg substrates were built using the EOS-nominal AlSi10Mg parameters (Table 2.1) on the EOS M280. To enable measurement of the melt pools (Section 2.2.3), the substrates were built with melt tracks oriented exclusively perpendicular to the beam travel direction used for the OLSB and OLP experiments themselves and the standard beam track rotation [41] was disabled. Additionally, the substrates were first surfaced using a facing mill before being placed back inside the EOS M280 for the OLSB and OLP experiments.







Figure 2.1: An example of a sampleFigure 2.2: An example of a sampleFigure 2.3: An example of a samplefrom the OLSB experiments.from the OLP experiments.from the MLP experiments.

The results from the OLSB experiments are reported by Narra [12, Ch. 5] and are only presented in this document in the form of Figure 2.13. A total of 24 different beam power and beam travel velocity combinations spanning the EOS L-PBF process space were chosen by Dr.

⁹ Zero Layer Single Bead (OLSB) experiments, i.e. zero layers of powder, single bead exposures.

¹⁰ Zero Layer Pad (OLP) experiments, i.e. zero layers of powder, multiple single bead exposures immediately adjacent to each other in a pad.

¹¹ Multi-Layer Pad (MLP) experiments, i.e. multiple layers of powder, used to build a solid block or pad.

Sneha Prabha Narra of CMU for the 0LP and MLP experiments. The hatch spacing (distance between the melt tracks as shown in Figure 1.7) was adjusted for the 0LP and MLP experiments such that the overlap (Figure 1.7) between the melt tracks (as a percentage of the melt pool width) remained equal to or greater than 7% – the expected overlap for the EOS-nominal parameters based on the 0LSB experiment results [12, Ch. 5]. Table 2.1 lists the process parameter combinations for the 0LP and MLP experiments. Both the 0LP and MLP experiments were run at a chamber preheat of 35 °C and a nominal¹² laser beam diameter of 100 μ m [75]. The MLP samples were built with a nominal layer thickness of 30 μ m, the EOS standard laser scan pattern rotation of 67° every layer [41] was used, but no contour¹³ beam passes were used.

¹² No direct, independent, measurements of the laser beam spot size were performed.

¹³ A "contour" refers to a laser beam pass which follows the perimeters (both internal and external) of a part. The contour passes are typically performed with different processing parameters than the bulk melt tracks.

Table 2.1: Process parameter combinations used for each OLP and MLP sample on the EOS M280 L-PBF machine as well as the EOS-nominal parameter combination.

Sample	Beam	Beam Velocity	Hatch
Number	Power (W)	(mm/s)	Spacing (µm)
EOS Nominal	370	1300	190
1	100	200	60
2	100	400	60
3	100	600	60
4	100	800	60
5	100	1000	60
6	200	200	130
7	200	400	120
8	200	600	110
9	200	800	100
10	200	1000	100
11	200	1200	100
12	200	1400	90
13	300	400	310
14	300	600	260
15	300	800	240
16	300	1000	130
17	300	1200	150
18	300	1400	100
19	370	400	350
20	370	600	300
21	370	800	270
22	370	1000	230
23	370	1200	200
24	370	1400	150

2.2.2 Sample Preparation

No post-build heat treatment was performed. The samples were sectioned (perpendicular to the beam travel direction) using a Wire EDM (Electrical Discharge Machine) at the Arconic facility. The EDM cutting process produces a relatively narrow heat-affected zone, reducing the influence of sectioning on measurements of the melt pools [76]. The sectioned samples were then hot-mounted, with the cut face visible, in Buehler[®] Konductomet sample pucks. After mounting, each sample was ground and polished according to ASTM E3-11, Table 7 [77]. To improve the visibility of the melt pool boundaries, the polished OLP samples were etched using Barker's agent (2% aqueous fluoroboric acid HBF₄) for 45 to 60 seconds as described by Fulcher et al. [62]. Finally, each polished sample was imaged using an Alicona Infinite-Focus optical microscope at an appropriate magnification. The OLP magnifications are listed in Table 2.2 while all MLP samples were imaged at 5x magnification.

2.2.3 OLP Measurement Techniques

The OLP melt pool widths, depths, and cross-sectional areas (Figure 1.7) were manually measured using the Image J software package [78]; Figure 2.4 shows an example micrograph from the OLP experiments. Due to the proximity of the melt tracks to each other, in practice, "half-widths" and "half-areas" were measured, as shown in Figure 2.4. Between 12 and 35 melt tracks were measured for each of the 24 samples (Table 2.1), with generally fewer melt tracks available for larger melt pools due to space restrictions on the AlSi10Mg build substrates (Section 2.2.1). A selection of OLP micrographs and the tabulated melt pool dimension measurements are provided in Appendix A.



Figure 2.4: A representative OLP micrograph, specifically from Sample #23 (370 W, 1200 mm/s). Note the half-width and depth measurement notations; the cross-sectional half-area is the region enclosed by the dotted white polygon. The coloration difference between the lower and upper halves of the micrograph demarcates the boundary between the additively manufactured AlSi10Mg substrate and the Al5083 base plate.

Due to the close spatial proximity of the melt tracks, the effect of residual heating (from previously deposited adjacent melt tracks) on the melt pool geometry was a concern. For this reason, the measured melt pool dimensions were plotted as a function of distance across the build surfaces; no trends¹⁴ were observed, suggesting that residual heating did not influence the melt pool geometry. Note that the relatively high thermal conductivity of AlSi10Mg likely makes this alloy less sensitive to residual heating than other alloys such as Ti-6Al-4V. Additionally, no measurements were included from melt pools closer than three melt pool widths from the edges¹⁵ of the build substrates. To quantify the uncertainty in the manual measurement of melt pool dimensions, a representative melt pool at each magnification was measured 10 times consecutively. The results of the measurement error and other measurement errors range from sample¹⁶ standard deviations that are 0.49% to 7.6% of the corresponding mean melt pool dimension.

¹⁴ If the residual heating from adjacent melt tracks was significant, it is expected that melt pools deposited last on the build substrates would be larger than melt pools deposited first, for the same processing parameters [184]. ¹⁵ The thermal conditions at the edges of the build substrates are expected to be different than those within in the

bulk region due to the extremely low thermal conductivity of metal powder relative to fused material [100], [210]. ¹⁶ In this context "sample" refers to the statistical term "sample of the population" and not the additively-produced 0LSB, 0LP, or MLP samples.

Magnification	Corresponding Sample Numbers	Std. Dev. of Half- Width Measurements (μm, % of mean)	Std. Dev. of Depth Measurements (μm, % of mean)	Std. Dev. of Half- Area Measurements (mm ² , % of mean)
10x	13 ¹⁷ , 14 ¹⁷ , 15 ¹⁷ , 19 ¹⁷ , 20 ¹⁷ , 21 ¹⁷ , 22, 23	7.0, 4.1	1.4, 0.49	1.8×10 ⁻³ , 5.5
20x	6 ¹⁸ , 7 ¹⁹ , 8, 9, 10, 11, 12, 16, 17, 18, 24	5.3, 7.4	0.67, 1.2	1.1×10 ⁻⁴ , 3.9
50x	1, 2, 3, 4, 5	2.2, 6.0	0.33, 1.7	3.9×10 ⁻⁵ , 7.6

Table 2.2: Corresponding measurement errors for the OLP experiments.

2.2.4 MLP Measurement Techniques

Bulk porosity was evaluated by first binarizing the optical micrographs (Figure 2.5) taken at 5x magnification (Section 2.2.2); where the binarization threshold was determined upon inspection of a bimodal intensity histogram. The bulk region, i.e. a zone away from the edges, (Figure 2.6) was then selected on the binary micrographs and a custom MATLAB script was used to identify and classify all of the pores within the selected region. A selection of MLP micrographs is available in Appendix B.

¹⁷ The melt pools in these samples penetrated deeper than (or close to) the height of the ASi10Mg build substrates, i.e. into the Al5083 build plate. Because of the differing thermal properties between these two aluminum alloys, the measured melt pool dimensions may not be wholly representative of those expected for an exclusively AlSi10Mg substrate. For validation of the presented results, see the high-degree of agreement between the 0LSB (for which there are no penetration concerns) melt pool widths and the 0LP melt pool widths in Figure 2.13.

¹⁸ Seven (7) out of 27 melt pools from Sample #6 were discarded as outliers. All of the outliers were approximately twice the width of the non-outlier melt pools. It is strongly suspected that incorrect processing parameters were programmed for those melt tracks.

¹⁹ One (1) out of 31 melt pools from Sample #7 was discarded as an outlier.





Figure 2.5: A representative MLP micrograph, specifically from Sample #24 (370 W, 1400 mm/s). The build direction (*z*-axis) is oriented upwards.

Figure 2.6: An example of a selected bulk region, indicated by the gray rectangle overlaid on the binarized version of the micrograph shown in Figure 2.5. The build direction (*z*-axis) is oriented upwards.

Pores are grouped into two categories: Pores greater than 40 μ m in effective circular diameter²⁰ and at least 90% circular²¹ are considered to be the result of "keyholing-mode" melt pools [74]. Pores greater than 40 μ m in effective circular diameter but less than 90% circular are considered "lack-of-fusion" flaws [74]. The described porosity binning thresholds are based on prior internal work and the work of Cunningham et al. [48], [79]. Pores with an effective circular diameter of less than 40 μ m were not considered. This method also discussed in Section 3.2.4 and the pore formation mechanisms are further discussed, in context, with the results in Section 2.3.4.

²⁰ The effective circular diameter is defined as the diameter of a circular pore with the same cross-sectional area as detected pore.

²¹ A pore is considered circular if greater than 90% of the pore pixels lie on top of the effective circular pore centered at the centroid of the original pore.

As-built 2D edge roughness was evaluated by first binarizing the optical micrographs (Figure 2.7) taken at 5x magnification (Section 2.2.2); where the binarization threshold was determined upon inspection of a bimodal intensity histogram. After binarization, internal porosity was automatically removed to prevent it from influencing the roughness calculations. The edge regions (two for each MLP sample) were then selected (Figure 2.8) and a custom MATLAB script was used to calculate four roughness measures: *Ra* (arithmetic average of absolute values) (2.1), *Rq* (root mean squared) (2.2), *Rz* (maximum peak-to-valley difference) (2.3), and *Rsk* (skewness) (2.4). Note that a negative *Rsk* value indicates the presence of sharp valleys (into the surface) while a positive *Rsk* value indicates the presence of sharp peaks (out of the surface) [80]. This methodology is also presented in Section 3.2.4.



Figure 2.7: A representative MLP micrograph, specifically from Sample #20 (370 W, 600 mm/s). The build direction (*z*-axis) is oriented upwards.

Figure 2.8: An example of a selected edge region, indicated by the gray rectangle overlaid on the binarized version of the micrograph shown in Figure 2.7. Note that the internal porosity has been automatically removed. The build direction (*z*-axis) is oriented upwards.

$$Ra = \frac{1}{n} \sum_{i=1}^{n} |y_i|$$
(2.1)

$$Rq = \sqrt{\frac{1}{n} \sum_{i=1}^{n} y_i^2}$$
(2.2)

$$Rz = \max_{i}(y_i) - \min_{i}(y_i)$$
(2.3)

$$Rsk = \frac{1}{n(Rq)^3} \sum_{i=1}^{n} y_i^3$$
(2.4)

Where *y* is the distance of a protuberance from the "mean" edge, and *n* is the number of protuberances along the edge.
2.3 Results

2.3.1 Melt Pool Geometry

Application of the process mapping approach to the OLP samples allowed for the generation of curves of constant melt pool geometries. Specifically, 2D linear interpolation was used to generate a dense matrix of melt pool geometry values (e.g. melt pool width) across beam power and beam travel velocity process space. This dense matrix was then queried such that a set of points in process space was produced at which the relevant melt pool geometry is the same. A smooth curve was then fitted to this set of points; in this case an exponential function of the form given in (2.5) is used. While it is common to fit linear curves to process maps in L-PBF and EB-PBF processes [29], the high thermal conductivity of AlSi10Mg shifts the process map into a heat transfer regime more commonly observed in the low power, low velocity LENS process (Figure 1.2) [31]. As such, the author felt that an exponential function was a more accurate representation of the underlying physics²², as well as providing better agreement with the data (see Table 2.3). A custom MATLAB script was used to automate the process described above.

²² For materials with a high thermal diffusivity (α), such as AlSi10Mg, typical L-PBF process space may include a region where the laser beam travel velocity is slow enough to be comparable to the "speed" at which heat can diffuse away from the melt pool and into the surrounding bulk material. In this region, linear increases in beam power do not result in a linear increase in melt pool size, instead the relationship is better described by an exponential function. This region exists for lower α materials (e.g. Ti-6Al-4V) in the LENS process owing to its lower laser beam velocity (Figure 1.2). The rate of thermal diffusion is described by the Fourier Number Fo = $\alpha t/L_c^2$ where t is a measure of time and L_c is a spatial measure describing the relevant geometry. The author directs the curious reader to the treatment of the Fourier Number, and transient heat conduction in general, provided by Bergman et al. [226, Ch. 5] for additional details.

$P = ab^{v}$

Where *P* is the beam power, *v* is the beam travel velocity, and *a* and *b* are the fitting parameters. Table 2.3: Comparison between power and linear fits of AlSi10Mg process map data.

Measurement	R² Value for Exponential Fit²³	R ² Value for Linear Fit ²³
Width	0.91	0.87
Depth	0.90	0.87
Area	0.87	0.80
Aspect Ratio	0.95	0.94

Figures 2.9 – 2.12 present process maps, respectively, for cross-sectional melt pool width, depth, area, and aspect ratio. The aspect ratio is defined as the depth divided by the half-width, e.g. an aspect ratio of 1.0 indicates a perfectly semicircular melt pool, an aspect ratio less than 1.0 indicates a shallow melt pool, and an aspect ratio greater than 1.0 indicates a deep and narrow melt pool. As expected, the process maps show that higher beam powers and lower beam velocities produce larger melt pools while lower beam powers and higher beam velocities result in smaller melt pools.

Figure 2.13 overlays the cross-sectional width results with the OLSB above-view (i.e. along the length of the of the melt track) width results reported by Narra [12, p. 96]. Note that relatively good agreement between the melt pool width measurement results is observed. Confirmation of such agreement may be useful to researchers in the future because above-view widths are often easier to measure than cross-sectional widths (as sectioning and polishing are not required) and above-view widths can be observed with various in-situ monitoring techniques [81] (see Chapter 7).

²³ The reported R² values are calculated based the agreement between the models (e.g. exponential or linear fits) and the data from each 0LP sample with a measurement (e.g. melt pool width) within the range presented in the corresponding process map. For example, the reported R² values for the melt pool width are based on data from the samples with a measured melt pool width between 100 μ m and 250 μ m (see the legend of Figure 2.9).





Figure 2.9: Process map of the cross-sectional melt pool **width**, developed from the OLP experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend.

Figure 2.10: Process map of the cross-sectional melt pool **depth**, developed from the OLP experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend.



Figure 2.11: Process map of the cross-sectional melt pool **area**, developed from the OLP experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend.



Figure 2.12: Process map of the cross-sectional melt pool **aspect ratio**, developed from the OLP experiment data. For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the left-right order shown in the legend.



Figure 2.13: A comparison between the OLSB above-view melt pool **widths** reported by Narra [12, p. 96] (dashed lines) and the cross-sectional (C-S) OLP melt pool **widths** reported in Figure 2.9 (solid lines). For the reader viewing the figure in grayscale, the top-bottom order of the lines of constant geometry matches the top-bottom order shown in the legend.

2.3.2 Distribution of Melt Pool Geometries

As a first step toward understanding the variability of melt pool geometry across process space, the measured size distributions (cross-sectional melt pool width, depth, and area) are shown as cumulative probability plots in Figures 2.14, 2.17, and 2.20. Normalization of the distribution curves was implemented by converting each individual measurement to its percent difference from the mean value for that power-velocity combination. Note the outlying measurements (first mentioned in Section 2.2.3) for Sample #6 and the single outlier for Sample #7; these measurements are not included in the analyses presented in Sections 2.3.1 and 2.3.3.

Normal probability plots²⁴ are shown in Figures 2.15, 2.16, 2.18, 2.19, 2.21, and 2.22 for each geometry measure (melt pool width, depth, and cross-sectional area) at both the process parameter combination closest to the EOS nominal process parameter combination and the process parameter combination which deviated the most from a normal distribution (excepting Samples #6 and #7). It is evident from both the cumulative probability and normal probability plots that the melt pools which deviate most significantly from the normal distribution form an upper tail. That is, while most of the melt pools follow a normal distribution, in some OLP samples several melt pools of a significantly larger size are present. The implications of this observation are discussed further in Section 2.4, and compared to the observed behavior of the Inconel 718 material system in Section 6.3.6.

To provide context for the use of confidence intervals based on Student's t-distribution [82, p. 419] in the previous section, the melt pool geometry data were quantitatively compared to their equivalent normal distribution. This comparison is shown graphically as the normal probability plots mentioned previously. The majority of the 0LP samples did not provide a sufficient number of measurements to perform a proper Chi-square (χ^2) test²⁵ [82, Ch. 10]; as a result, Table 2.4 and the legends of Figures 2.14, 2.17, and 2.20 instead report R^2 fit values between each 0LP data set and its equivalent normal distribution, both of which have been

²⁴ In a normal probability plot the data are sorted as they would be in a cumulative distribution plot and then they are plotted on a non-linear vertical axis representing the normal order statistic medians. If the data are samples which "come from a population with a normal distribution" [227] then they will fall along a straight line [227]. Note that in this context "sample" refers to the statistical term "sample of the population" and not the additively-produced OLSB, OLP, or MLP samples. Note also that in the implementation [227] used to generate the normal probability plots in this document, the equivalent normal distribution is calculated using only data from the second and third data quartiles.

 $^{^{25}}$ The standard rule of thumb is that 5 – 8 bins containing a minimum of 5 measurements (i.e. 25 – 40 measurements) are required to perform a valid Chi-square test [82, p. 307].

linearized. The majority of the cross-sectional geometry measurements follow normal distributions with the outliers following the trend discussed above. Overall, the agreements of the depth and area cumulative distributions with their equivalent normal distributions are weaker than observed for the cross-sectional width measurements.



Figure 2.14: Normalized cumulative probability plots of cross-sectional melt pool widths for all 24 0LP samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.



0.98 0.95 0.90 0.75 Probability 0.50 0.25 0.10 0.05 0.02 0 10 -10 -5 5 Percent Difference from the Mean

Figure 2.15: Normal probability plot of the measured **widths** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.

Figure 2.16: Normal probability plot of the measured widths for the process parameter combination showing the greatest deviation from a normal distribution (Sample #22).



Figure 2.17: Normalized cumulative probability plots of cross-sectional melt pool **depths** for all 24 0LP samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.





Figure 2.18: Normal probability plot of the measured **depths** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.

Figure 2.19: Normal probability plot of the measured **depths** for the process parameter combination showing the greatest deviation from a normal distribution (Sample #12).



Figure 2.20: Normalized cumulative probability plots of cross-sectional melt pool **areas** for all 24 0LP samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.



Figure 2.21: Normal probability plot of the measured **area** for the process parameter combination closest to the EOS nominal process parameter combination (Sample #24). Experimental points far away from the line indicate a deviation from a normal distribution.



Figure 2.22: Normal probability plot of the measured **areas** for the process parameter combination showing the greatest deviation from a normal distribution (Sample #12).

Table 2.4: Goodness-of-fit between the measured melt pool distributions and their equivalent normal distributions. R^2 values less than 0.80 are highlighted for convenience.

Sample	P^2 (width)	P^2 (donth)	P^2 (area)
Number		x (deptil)	n (alea)
1	0.95	0.90	0.85
2	0.94	0.93	0.90
3	0.91	0.92	0.96
4	0.96	0.11	0.17
5	0.87	0.86	0.79
6	0.65	0.59	-43
7	-2.8	-120	-84
8	0.95	0.81	0.94
9	0.93	0.92	0.76
10	0.95	0.81	0.86
11	0.88	0.74	0.94
12	0.90	-2.8	-1.7
13	0.94	0.67	0.91
14	0.96	0.96	0.97
15	0.80	0.92	0.95
16	0.86	0.86	0.84
17	0.93	0.81	0.82
18	0.81	-0.60	-1.1
19	0.91	0.98	0.84
20	0.84	0.88	0.93
21	0.92	0.87	0.92
22	0.42	0.94	0.94
23	0.97	0.83	0.95
24	0.46	0.66	0.70

2.3.3 Variability of Melt Pool Geometry across Process Space

The variability of melt pool geometry across process space was also investigated. Figures 2.23 - 2.25 show the standard deviation (as a percentage of the mean melt pool dimension) for, respectively, the OLP cross-sectional melt pool width, depth, and area. The measured variabilities in melt pool width, depth, and area respectively range from approximately 3.8% - 19%, 2.9% - 32%, and 3.5% - 44% of the mean melt pool dimension. With few exceptions, the melt pool dimension variabilities are as large as or larger than the measurement errors reported in Section 2.2.3.

Interestingly, no significant trend for this measure of variability presents itself, regardless of whether a function of beam power or beam velocity is considered. Note that there does appear to be a region of particularly high variability encompassing Samples #16 (300 W, 1000 mm/s) and #17 (300 W, 1200 mm/s), the cause of which is unknown. The aforementioned lack of a trend in melt pool variability across process space may be more easily visualized using 2D linearly-interpolate heat maps, such as those presented in Figures 2.26 – 2.28.



Figure 2.23: The variability (standard deviation) in the melt pool **widths** as a percentage of the mean width. Samples are grouped by beam velocity, with each beam power denoted by a different bar hue as shown in the legend. The error bars represent a 95% confidence interval about the sample²⁶ (percent) standard deviation.

²⁶ In this context the term "sample" refers to the statistical definition of a "sample of a population" and should not be confused with the OLSB, OLP, or MLP "samples."



Figure 2.24: The variability (standard deviation) in the melt pool **depths** as a percentage of the mean depth. Samples are grouped by beam velocity, with each beam power denoted by a different bar color as shown in the legend. The error bars represent a 95% confidence interval about the sample²⁶ (percent) standard deviation.



Figure 2.25: The variability (standard deviation) in the melt pool **areas** as a percentage of the mean area. Samples are grouped by beam velocity, with each beam power denoted by a different bar color as shown in the legend. The error bars represent a 95% confidence interval about the sample²⁶ (percent) standard deviation.



Figure 2.26: Interpolated heat map of the variability (standard deviation) in the melt pool **widths** as a percentage of the mean width. Note that geometric variability may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

Figure 2.27: Interpolated heat map of the variability (standard deviation) in the melt pool **depths** as a percentage of the mean depth. Note that geometric variability may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.



Figure 2.28: Interpolated heat map of the variability (standard deviation) in the melt pool **areas** as a percentage of the mean area. Note that geometric variability may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

2.3.4 Bulk Porosity in As-Built MLP Samples

The bulk porosity in the MLP samples was calculated and categorized as described in Section 2.2.4. Figure 2.29 is a 2D linearly interpolated heat map showing the porosity content attributed to the lack-of-fusion mechanism [74]. Because the hatch spacing was modified to ensure lateral overlap between the melt pools (Table 2.1), it is expected that any lack-of-fusion porosity is the result of the melt pools not penetrating sufficiently deep through the 30 μ m thick powder layer (i.e. under-melting). For reference, the fitted curves of 30 μ m and 40 μ m deep melt pools are overlaid on top of the heat map.

Tang et al. developed a model which predicts the regions of process space where lack-offusion porosity can be expected [72]. The process parameter combinations which are predicted to produce lack-of-fusion porosity are indicated by the larger open white circles in Figure 2.29. It is clear that there is relatively poor agreement between the predictions and the measured amounts of lack-of-fusion porosity. Interestingly, the model results are highly sensitive to the measured melt pool depths and to an even greater extent, the melt pool widths. The smaller open white circles in Figure 2.29 indicate the predicted porosity threshold assuming that the melt pools in the MLP samples are 75% deeper and 10% wider than the corresponding melt pools in the 0LP samples. Such an increase in melt pool depth is similar to that observed for Inconel 718 in Section 6.3.4 when melt pools exposed on top of a powder layer are compared to those exposed on top of a bare substrate. While the data in Section 6.3.4 do not demonstrate a similar increase in melt pool width, this may be material-dependent and only a small increase in width is required for the Tang model to accurately predict the measured lackof-fusion porosity threshold. Therefore, the author hypothesizes that a melt pool produced by a

45

given set of process parameters would be larger during MLP experiments, and therefore less likely to result in lack-of-fusion porosity, than during OLP experiments. It would be extremely worthwhile to perform 1LSB experiments (similar to those described in Section 6.2.1) for AlSi10Mg and compare them to the OLP results presented in this chapter to determine if the above hypothesis is correct (see Section 8.3). The effects of a powder layer on melt pool geometry are discussed in more detail in Section 6.3.4.



Figure 2.29: An interpolated heat map of the porosity content attributed to the **lack-of-fusion** mechanism. Note the overlaid lines of constant melt pool depth. The circles indicate process parameter combinations which are predicted to cause lack-of-fusion porosity based on the model proposed by Tang et al. [72] under two different sets of assumptions. Note that bulk porosity may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

Figure 2.30 is a 2D linearly interpolated heat map showing the porosity content attributed to the keyholing mechanism [74]. Keyhole-mode melting occurs in the high energy density (high beam power, low beam velocity) region of process space, where periodic vaporization of the molten material can occur. Keyhole-mode melting is associated with high aspect ratio melt pools [31], [83]. For reference, the fitted lines of melt pools with aspect ratios of 1.0 and 1.3 are overlaid on top of the heat map. Note that demarcating the region of process space sensitive to keyholing-induced porosity with a line of constant aspect ratio will likely produce a conservative estimate of usable process space at higher beam velocities. As is apparent in Figure 2.30, keyholing porosity is virtually non-existent for beam velocities higher than 800 mm/s, even for higher aspect ratio melt pools. Prior welding research [84] and a growing body of Additive Manufacturing research [85] indicate that regimes of stable keyhole-mode melt pools exist. In these regimes, the melt pool morphology is such that periodic collapse of the vapor pocket, and subsequent entrapment of the keyholing pore by the liquid front, is relatively infrequent.



Figure 2.30: An interpolated heat map of the porosity content attributed to the **keyholing** mechanism. Note the overlaid lines of constant melt pool aspect ratio. Note that bulk porosity may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

2.3.5 Edge Roughness of As-Built MLP Samples

Figures 2.31 – 2.34 show analyses of the edge roughness of 23 of the 24 MLP samples²⁷ in the form of 2D linearly interpolated heat maps. Each reported value represents an average of the values determined for the two vertical (build direction) edges of each MLP sample. Each measured edge was approximately 1 cm in length. For reference, the fitted lines of 30 μ m and 40 μ m deep melt pools are overlaid on top of the heat map. A region of relatively high edge roughness is evident in the low energy density regime of process space. In particular, the high edge roughness values for the *Ra*, *Rq*, and *Rz* metrics appear limited to the lack-of-fusion region (introduced in Section 2.3.4), where the melt pools are shallower than the 30 μ m layer thickness.

Notably, negative *Rsk* values are not limited to the lack-of-fusion region, indicating that sharp valleys on the surface may be caused by mechanisms in addition to surface-connected lack-of-fusion porosity. Internal modeling work by Christopher Kantzos of CMU has demonstrated that surface valleys are a significant source of stress concentrations (potential crack initiation sites) under certain tensile loading conditions. It can therefore be inferred that a highly negative *Rsk* value may indicate that a part has a decreased resistance to fatigue failure. Across process space, the *Ra* metric ranges from 22 µm to 130 µm, the *Rq* metric ranges from 29 µm to 170 µm, the *Rz* metric ranges from 110 µm to 510 µm, and the *Rsk* metric ranges from -0.44 µm to 0.35 µm.

²⁷ The edges of MLP Sample #12 were not available for measurement (see Appendix B).



Figure 2.31: An interpolated heat map of the **Ra** edge roughness measure. Note the overlaid lines of constant melt pool depth. Note that edge roughness may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.



Figure 2.33: An interpolated heat map of the **Rz** edge roughness measure. Note the overlaid lines of constant melt pool depth. Note that edge roughness may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.



Figure 2.32: An interpolated heat map of the *Rq* edge roughness measure. Note the overlaid lines of constant melt pool depth. Note that edge roughness may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.



Figure 2.34: An interpolated heat map of the *Rsk* edge roughness measure. Note the overlaid lines of constant melt pool depth. Note that edge roughness may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

2.3.6 A Robust Processing Window

As discussed in Section 1.2, the primary goal of process mapping is to provide insight into the correlation between process parameters and process outcomes. Because bulk porosity is a major factor in a part's final quality, particularly as it relates to fatigue life [79], [86], Figure 2.35 presents a desirable processing window inside of which porosity due to lack-of-fusion and keyholing can be largely avoided when building the bulk region of a part. This processing window is bounded by a curve of minimum melt pool depth and a curve of maximum melt pool aspect ratio. Four representative MLP micrographs from across process space are also included for reference. This process map is an extension of the process map presented by Narra [12, p. 102], further informed by a quantitative analysis of the bulk porosity content and the higher degree of confidence in the cross-sectional melt pool measurements. Note that this process window is contingent upon the controlled parameters such as hatch spacing (Table 2.1), powder layer thickness (30 µm), and build chamber temperature (35 °C).



Figure 2.35: A proposed processing window for AlSi10Mg, assuming that hatch spacing is controlled to ensure lateral overlap and 30 μ m thick powder layers are used. Parts built with parameters located above the 40 μ m melt pool depth curve and below the 1.3 aspect ratio curve are expected to have low levels of porosity due to lack-of-fusion and keyholing. Four representative MLP micrographs are also shown. The black cross-mark at 370 W, 1200 mm/s denotes the default EOS M290 process parameters for AlSi10Mg.

2.4 Discussion and Summary

In this chapter, correlations between process parameters (beam power and beam travel velocity) and cross-sectional melt pool geometry (width, depth, area, and aspect ratio) are presented in the form of process maps for the L-PBF-processed AlSi10Mg alloy. The process map developed for cross-sectional width is compared to the process map reported by Narra [12, Ch. 5] which was derived from above-view measurements of the melt tracks. Favorable agreement between these two maps was found suggesting that melt pool widths collected via certain in-situ process monitoring techniques could accurately represent melt pool dimensions which are more commonly measured through ex-situ destructive testing.

While process maps of L-PBF-processed AlSi10Mg cross-sectional melt pool geometry have been reported by Narra [12, Ch. 5], they were based on single measurements of each process parameter combination. The process maps presented in this chapter are based on data from multiple cross-sections at each process parameter combination – allowing for the presentation of process maps with an increased, and quantifiable, level of confidence. Additionally, the relatively large number of cross-sectional measurements allowed for an investigation of the statistical distribution of melt pool dimensions for each process parameter combination. Analysis of the distributions reveals that cross-sectional melt pool widths, depths, and areas primarily follow a normal distribution with the exception of a handful of outliers (at certain process parameter combinations) which clearly diverge from a normal distribution. Interestingly, the divergent melt pools almost exclusively formed an upper tail, that is, the divergent melt pools were significantly larger than the majority of the melt pools produced by that process parameter combination. The large sample size also allowed for the quantification of the variability of the melt pool dimensions across process space – critical information for designers as they work at the edges of viable L-PBF processing space. Notably, no correlation between variability and processing parameters was observed.

The bulk porosity contents of L-PBF-processed AlSi10Mg samples are reported across process space. Pore categorization based on size and morphology allowed for the identification of regions of process space dominated by lack-of-fusion and keyholing flaw formation mechanisms. The experimentally-determined lack-of-fusion region was compared to the region predicted by the model proposed by Tang et al. [72]; the existence of a notable discrepancy suggests that 0LP melt pool geometry may not be fully descriptive of MLP melt pool geometry

52

due to the effects of the powder layer. The region of process space where keyholing porosity was observed is bounded by a curve of constant melt pool aspect ratio, although this boundary is likely conservative at higher beam velocities as discussed previously. Four 2D edge roughness measures were used to evaluate the as-built surfaces of the MLP samples; unsurprisingly, samples within the lack-of-fusion regime exhibited noticeably rougher surfaces although highly negative values of the skewness metric (*Rsk*) were also observed in regions outside of the lack-of-fusion regime. It should be noted that while informative, the reported edge roughness values and trends are not necessarily indicative of the values and trends which would be observed if the MLP samples had (as is standard in most L-PBF systems) been built with contours instead of only bulk processing parameters and bulk laser scan strategies.

Finally, the quantitative porosity results are combined to generate an L-PBF processing window that provides a range of processing parameters expected to produce AlSi10Mg parts with minimal bulk porosity due to lack-of-fusion or keyholing. This processing window may be treated as an extension of the window proposed by Narra [12, Fig. 5.24] with an increased level of fidelity. Given the lack of keyholing porosity at higher beam velocities, it is likely that AlSi10Mg parts could be built with higher beam powers and travel velocities than is currently standard for the EOS M90 without a detrimental decrease in part quality. Such a change would be advantageous as increases in beam power and velocity allow for an increased material deposition rate [33]. Note, however, that while neither an increase in melt pool variability nor a

change in melt pool morphology were observed at the higher beam velocities tested²⁸, a surface tension instability phenomenon known as "balling" is generally known to occur at high beam velocities [87], [88], particularly in the presence of a powder layer [69]. Balling is discussed extensively throughout Chapter 6 and the behavior of AlSi10Mg melt pools is compared to the behavior of Inconel 718 melt pools in Section 6.3.6.

²⁸ More specifically, the onset of balling occurs as the length to width ratio of the melt pool increases, i.e. as the melt pool gets narrower. Because of the high thermal diffusivity of AlSi10Mg, the melt pools tend to remain short relative to their width over the region of process space explored in this thesis.

3 Topic 2: The Effect of Non-Standard Metal Powders on L-PBF Part and Process Quality

3.1 Background and Literature Review

One of the primary impediments to the wide-spread adoption of additive technologies is the current cost and sourcing restrictions of the feedstock for PBF processes [3], [4]. Currently, machine manufactures strongly recommend that their customers use only powders from vetted vendors, produced via specific processes, and with restrictive size distributions [51]. These restrictions are intended to ensure production of parts with the quality and properties guaranteed by the machine manufacturers, as well as to prevent damage to the machines themselves.

The work presented in this chapter seeks to demonstrate that a relatively wide range of metal powder types (e.g. size distributions and production methods) can be used to successfully build parts in an L-PBF system. This work specifically targets the use of larger-than-standard powders in L-PBF because the cost of a powder lot generally decreases (for a given material) as the Powder Size Distribution (PSD) widens to include the larger particles produced during most traditional powder production techniques. For example, less than 15% of a powder batch produced via gas atomization would be considered usable by L-PBF processes if a traditional size cutoff of 60 µm is used; but if powder particles up to 120 µm in diameter become acceptable for L-PBF, that fraction increases to almost 40% [89, Ch. 5].

Most powders used in metal AM processes are produced by either gas or plasma atomization [51]; owing to the methods' relatively low cost and ability to produce spherical powders [51]. Plasma Rotating Electrode Powder (PREP) is more expensive to produce, but is

55

generally considered to be of high purity and extremely spherical relative to other production methods [90]. Hydride-dehydride (HDH) is a powder production technique unique to titanium alloys; while it is substantially cheaper than either atomization or PREP, the powders tend to be irregular in shape [90], [91]. A degree of sphericity can be achieved with HDH powder if Plasma Spherodization (P-S) is performed as an additional step [90], [91].

The characteristics of the powder used can affect the ability of the recoater blade to spread the powder, the packing density of the powder [92], and the bulk porosity and surface roughness of the final parts. While there have been several studies investigating the spreadability and flow characteristics of AM powders, particularly from a Discrete Element Modeling (DEM) approach [93]–[95], the author has found only a limited body of work linking PSDs to final part quality. Specifically, Spierings et al. [96] investigated the influence of three different stainless steel 316L powders on surface roughness, bulk porosity, and tensile strength, finding variation across the Powder Systems but no clear trend with respect to the PSD. The transfer of porosity from powder particles to final parts was investigated by Cunningham et al. [48] in the L-PBF process. Saint John et al. [97] performed mechanical testing on components built with four different Inconel 718 powders and finding limited variation across the Powder Systems with the exception of the measured impact energy. Finally, Gu et al. [98] performed mechanical testing on components built with three different Ti-6Al-4V powders for which no significant differences in mechanical properties were found. The alternate Powder Systems investigated by Spierings et al. and Gu et al. did not exceed 50 µm in diameter; while the largest powder in the study performed by Saint John et al. had a mean particle diameter of 47 μ m. In contrast, the Powder Systems studied in this work contain particles up to 120 μ m in diameter (Table 3.2).

In this chapter, the effects non-standard Powder Systems (PS) on build and as-built part quality are explored. Specifically, three non-standard Ti-6Al-4V (Ti64) Powder Systems and one non-standard Inconel 718 (In718) Powder System are successfully used to build a standard set of test artifacts. The performance of these four non-standard Powder Systems is compared to that of the EOS nominal Ti64 Powder System. Qualitative observations of build quality as well as a quantitative in-situ analysis of powder spreading anomalies (using the methodology described in Chapter 4) are presented. Build quality is quantified by measurements of bulk porosity content and the 2D edge roughness of the as-built parts. The work presented in this chapter was supported by the America Makes project: "A Database Relating Powder Properties to Process Outcomes for Direct Metal Additive Manufacturing" (project number 4028.001).

3.2 Experimental Design and Methods

3.2.1 Feedstock

Standardized test artifacts were built with four non-standard and one standard (nominal) Powder Systems. The nominal characteristics of these powders are summarized in Table 3.1.

Table 3.1: Nominal Powder System characteristics as reported by the powder manufacturers.

Powder System ²⁹	PS #1	PS #2	PS #3	PS #4	PS #5
Manufacturer	EOS Standard	А	В	В	С
Material	Ti64	Ti64	Ti64	Ti64	In718
Production Process	Gas Atomized	PREP	HDH + P-S	HDH + P-S	Gas Atomized
Mesh Size	-230	-170	-200/+325	-140/+325	-170/+800
Particle Diameter (µm)	< 63	< 88	74 – 44	105 – 44	< 88

In contrast, Table 3.2 shows volume-weighted Powder System characteristics as measured independently by collaborators. Specifically, the data reported for PS #1 were collected and analyzed by Ross Cunningham and Prof. Anthony Rollett of CMU while the data reported for PS #2 – #5 were collected and analyzed by Hengfeng Gu and Prof. Ola Harrysson of North Carolina State University (NCSU). Table 3.2 is included in this thesis solely to provide background to the readers and allow for the presentation of results in the context of the measured PSDs.

Table 3.2. Volume-weighted Powde	System characteristics as measured	by collaborators at CMU and NCSU
Table 3.2. Volume-weighteu rowue	System characteristics as measured	by collaborators at civio and incoo.

	PS #1	PS #2	PS #3	PS #4	PS #5
Lower 10% Threshold 30 (µm)	20	50	60	35	15
Mean (µm)	34	78	91	62	58
Upper 10% Threshold ³¹ (µm)	50	105	120	90	105

²⁹If the reader wishes to compare any results reported in this thesis to the America Makes report [53], be advised that the Powder System numbering scheme is not consistent. Please refer to the following conversion key: thesis (America Makes), PS #1 (PS #8), PS #2 (PS #3), PS #3 (PS #4), PS #4 (PS #5), and PS #5 (PS #7).

³⁰ Approximately 10% of the powder particles (by volume) in the sample population have sizes smaller than the stated value, to the nearest available bin size.

³¹ Approximately 10% of the powder particles (by volume) in the sample population have sizes larger than the stated value, to the nearest available bin size.

3.2.2 Build Conditions

A CAD (Computer Aided Design) rendering of the test artifacts are shown in Figure 3.1. Each build consists of a single NIST part [99] located in the center of the build area and surrounded by eight vertical cylinders, each 2 cm tall and 2 cm in diameter. The builds were performed on an EOS M290 L-PBF machine at CMU's NextManufacturing Center; no modifications (mechanical or software) were made to the powder spreading and handling mechanisms within the EOS M290. For each build, the test artifacts were built in the same orientation and location in the build area. Solid supports, 4.5 mm tall, were used for all builds; with a single layer skipped at the support-part interface.





The process parameters used for the five builds are listed in Table 3.3 and were chosen to enable proper spreading of the powders as well as mitigate the occurrence of bulk porosity. Specifically, the nominal powder layer thickness was first chosen such that the effective layer thickness would be greater than the diameter of the largest powder particles reported by the powder manufacturer³². Note that the effective layer thickness is larger than the nominal layer

 $^{^{32}}$ Administrative restrictions prevented independent measurement of the PSDs (Table 3.2) prior to the start of the PS #1 – PS #5 builds.

thickness due to layer-wise post-fusion consolidation of the powder. Based on the work of Jacob et al. [92], a consolidation factor³³ (\varkappa) of 0.40 (40%) was chosen. Under these assumptions, nominal powder layer thicknesses of 30 µm and 60 µm result in effective powder layer thicknesses of 50 µm and 100 µm, respectively. Based on these predictions, a nominal powder layer thickness of 60 µm was considered sufficient for PS #1 – #5. An extensive discussion of layer-wise post-fusion powder consolidation is available in Section 6.2.1.

	PS #1	PS #2	PS #3	PS #4	PS #5
Build Chamber Temperature (°C)	35	35	35	35	80
Nominal Beam Diameter ³⁴ (μm)	100	100	100	100	100
Starting Layer Thickness (µm)	60	60	60	60	60
Nominal Layer Thickness ³⁵ (µm)	60	60	60	60	60
Raster (bulk) Beam Power (W)	340	340	340	340	370
Raster (bulk) Beam Velocity (mm/s)	1250	1250	1250	1250	1000
Raster (bulk) Hatch Spacing (µm)	120	120	120	120	120
Raster (bulk) Stripe Width (mm)	5	5	5	5	10
Inner Post-Contour Beam Power (W)	190	190	190	190	270
Inner Post-Contour Beam Velocity (mm/s)	1200	1200	1200	1200	800
Outer Post-Contour Beam Power (W)	190	190	190	190	270
Outer Post-Contour Beam Velocity (mm/s)	1250	1250	1250	1250	800
Post-Contour Offset ³⁶ (µm)	30	30	30	30	12

Table 3.3: Processing parameters used for each build on the EOS M290 L-PBF machine.

³³ Where the "consolidation factor" is defined as: $\kappa = 1 - \rho_{powder} / \rho_{fused}$, where ρ is the density of the material.

 $^{^{34}}$ The D86 beam diameter was measured to be approximately 90 μm during the machine maintenance temporally closest to the builds enumerated in Table 3.3.

 $^{^{35}}$ Note that this is the nominal powder layer thickness; layer-wise post-fusion consolidation of the powder results in an effective powder layer thickness that is greater than the nominal value. For an assumed consolidation percentage of 40%, the effective layer thickness will be approximately 100 μ m. In other words, the effective powder layer thickness is greater than the mean particle size for all of the Powder Systems. See Section 6.2.1 for a more detailed discussion of this topic.

³⁶ This is the separation distance (i.e. hatch spacing) between the two post-contour melt tracks.

EOS nominal parameters are available for 60 µm layers for the Ti64 material system and were used for the PS #1 – #4 builds. In contrast, the largest available EOS layer thickness for In718 is 40 μ m, therefore custom process parameters were designed by the author. First, the beam power and beam travel velocity were chosen such that the melt pools in the bulk region would fully penetrate through the nominal layer thickness. This choice was informed by an experimental process map developed for the In718 material system in the EOS M290 machine by Narra [12, Ch. 6]. This process map is based on cross-sectional measurements of individual melt tracks exposed on a bare build substrate devoid of powder. A simplified version of this melt pool cross-sectional depth process map is shown in Figure 3.2. The chosen laser beam power (370 W) and travel velocity (1000 mm/s) result in an expected melt pool depth of approximately 100 µm. In order to ensure lateral overlap between adjacent melt tracks (Figure 1.7), a conservative hatch spacing of 120 µm was chosen. This hatch spacing results in an expected overlap equal to 17% of the expected melt pool cross-sectional width (180 µm based on the experimental process maps reported in [12]). The EOS nominal stripe width of 10 mm was maintained. Finally, the contour process parameters were chosen to produce melt pools with a reduced depth (70 μ m) as the differing thermal conditions near the edge of a part [100] are expected to increase the melt pool depth relative to its depth within a bulk region; this behavior is discussed further in Chapter 7. Note that as discussed in Section 6.3.4, the use of a process map based on experiments performed on a bare substrate likely resulted in conservative estimates of melt pool size in the presence of a powder layer. Additional details regarding desirable L-PBF processing windows and undesirable defects are available in Chapters 2 and 6.



Figure 3.2: An experimental process map of melt pool cross-sectional depth for the In718 material system in the EOS M290 machine. This process map is based on data reported by Narra [12, Ch. 6]

3.2.3 Sample Preparation

All of the builds were stress-relieved (annealed) using the EOS-recommended furnace profiles, which are reported in Table 3.4. After the stress-relief, the test artifacts were removed from the build plate using a gravity-fed band saw.

	Ti64 [101]	In718 [102]
Ramp Up (°C/min)	+20	+20
Soak Temperature (°C)	650	1065
Soak Time (min)	180	60
Nominal ³⁷ Ramp Down (°C/min)	-20	-20
Gas Environment	argon	argon

 $^{^{37}}$ This was the target cool down rate; the achieved (i.e. actual) cool down rates ranged between –5 °C/min and –1 °C/min.

After removal from the EOS build plate, five of the eight cylinders from each build were sectioned along their build (*z*-axis) using either a low-speed saw with diamond wafering blade³⁸ or a Wire EDM (Electrical Discharge Machine). One half of each of those cylinders was then hot-mounted, with the cut face visible, in Buehler[®] Konductomet sample pucks. After mounting, each sample was ground and polished according to ASTM E3-11, Table 6 [77]. During preparation of the Ti64 samples, special care was taken to use 0.06 μ m colloidal silica (Buehler[®] MasterMetTM) mixed with hydrogen peroxide (H₂O₂) in a 5:1 ratio as the lubricant for all of the polishing steps, per a procedure developed by Dr. Jason Fox of CMU. Finally, each polished sample was imaged using an Alicona Infinite-Focus optical microscope at 5x magnification.

3.2.4 Cylindrical Sample Measurement Techniques

Bulk porosity was evaluated by first binarizing the optical micrographs taken at 5x magnification (Figure 3.3); where the binarization threshold was determined upon inspection of a bimodal intensity histogram. The bulk region, i.e. a zone away from the edges, (Figure 3.4) was then selected on the binary micrographs and a custom MATLAB script was used to identify and classify all of the pores within the selected region. A selection of cylindrical sample micrographs is available in Appendix C.

³⁸ It was later learned that abrasive blades (as opposed to diamond wafering blades) are recommended by Struers Inc. for sectioning Ti64 and In718 samples in order to decrease the cutting time and reduce the wear on the blade.



Figure 3.3: Original micrograph of a cylinder built with PS #2. The build direction (*z*-axis) is oriented upwards.



Figure 3.4: An example of a selected bulk region, indicated by the gray rectangle overlaid on the binarized version of the micrograph shown in Figure 3.3. The build direction (*z*-axis) is oriented upwards.

As discussed by Saltykov [103], the 2D pore size distribution underestimates the true, 3D pore sizes for a given sample. While approximations of 3D pore size based on 2D cross-sectional data can be performed [104], these analyses make a variety of assumptions regarding pore sphericity and the behavior of the pores during polishing (Section 3.2.3) that the author felt were not necessarily well-matched to the experimental data. Therefore pore classification was performed based on the 2D pore size distribution, as described below.

Pores are classified into four categories: Pores less than 20 μ m in effective circular diameter³⁹ are considered to have been natively present in the powder particles as a result of their production [74], [79]. Pores greater than 40 μ m in effective circular diameter and at least

³⁹ The effective circular diameter is defined as the diameter of a circular pore with the same cross-sectional area as detected pore.

90% circular⁴⁰ are considered to be the result of "keyholing-mode" melt pools [74]. Pores greater than 40 μ m in effective circular diameter but less than 90% circular are considered "lack-of-fusion" flaws [74]. Pores greater than 20 μ m but less than 40 μ m in effective circular diameter are left "unclassified," due the overlap in the sizes of the pores formed by the mechanisms enumerated above. The described porosity binning thresholds are based on prior internal work and the work of Cunningham et al. [48], [79]. Pores with an effective circular diameter of less than 5 μ m were not considered, given the resolving power⁴¹ of the microscope at the magnification used. Note that the classification of keyhole and lack-of-fusion pores is the same as in Section 2.2.4; further discussion of the lack-of-fusion and keyhole porosity formation mechanisms is available in Section 2.3.4.

As-built 2D edge roughness was evaluated by first binarizing the optical micrographs taken at 5x magnification (Figure 3.5); where the binarization threshold was determined upon inspection of a bimodal intensity histogram. After binarization, internal porosity was automatically removed to prevent it from influencing the roughness calculations. The edge regions (two for each cylindrical sample) were then selected (Figure 3.6) and a custom MATLAB script was used to calculate four roughness measures: *Ra* (arithmetic average of absolute values) (3.1), *Rq* (root mean squared) (3.2), *Rz* (maximum peak-to-valley difference) (3.3), and *Rsk* (skewness) (3.4). Note that a negative *Rsk* value indicates the presence of sharp valleys

⁴⁰ A pore is considered circular if greater than 90% of the pore pixels lie on top of the effective circular pore centered at the centroid of the pore.

 $^{^{41}}$ At 5x magnification, the resolving power was determined to be 3 μm based on visual inspection of a 1951 USAF resolution test target.

(into the surface) while a positive *Rsk* value indicates the presence of sharp peaks (out of the surface) [80]. This methodology also presented in Section 2.2.4.



Figure 3.5: Original micrograph of a cylinder built with PS #2. The build direction (*z*-axis) is oriented upwards.



Figure 3.6: An example of a selected edge region, indicated by the gray rectangle overlaid on the binarized version of the micrograph shown in Figure 3.5. Also note that the internal porosity has been automatically removed. The build direction (*z*-axis) is oriented upwards.

$$Ra = \frac{1}{n} \sum_{i=1}^{n} |y_i|$$
(3.1)
$$Rq = \sqrt{\frac{1}{n} \sum_{i=1}^{n} y_i^2}$$
(3.2)

$$Rz = \max_{i}(y_i) - \min_{i}(y_i)$$
(3.3)

$$Rsk = \frac{1}{n(Rq)^3} \sum_{i=1}^{n} y_i^3$$
(3.4)

Where *y* is the distance of a protuberance from the "mean" edge, and *n* is the number of protuberances along the edge.

3.3 Results

3.3.1 Qualitative Build Quality and Observations

The non-standard powder spreadability was qualitatively evaluated in a "go/no-go" context. All of the Powder Systems in Table 3.1 spread reliably throughout their respective builds. During the building process, anomalies in the powder bed were noticed for several nonstandard Powder Systems. These anomalies⁴², shown in Figure 3.7, took the form of clumps of dark particles at the edges of the stripes (Figure 1.8) during melting. These anomalies were most prominent in PS #2 and PS #3, with some similar clumping observed during the PS #4 build. No such anomalies were observed during either the PS #5 or PS #1 (EOS nominal) builds.



Figure 3.7: Images taken through the operator window of the EOS M290 during the building process.

(a): PS #1 (nominal), a desirable powder bed, no visible dark particles.

- (b): PS #2, note the ridges of large clumps of dark particles highlighted by the red oval.
- (c): PS #3, note the ridges of large clumps of dark particles.
- (d): PS #4, note the ridges of large clumps of dark particles.
- (e): PS #5, a desirable powder bed, no visible dark particles.

⁴² Note that in this context, "anomaly" is used as a general term for an observable disturbance on the powder bed and does not specifically refer to one of the powder bed anomalies enumerated in Chapter 4.
The test artifacts were successfully built with all of the tested Powder Systems (Figure 3.8). The discolorations visible in Figure 3.8 are a result of oxidation during the stress-relief process (this is mitigated by performing the heat treatment in an inert argon atmosphere, but it does not eliminate the oxidation completely). Normally, this oxidation layer would be easily removed by shot peening, but this has not been done in order to preserve the as-built surface finish for this study. The severe oxidation in Figure 3.8d is the result of an accidental loss of argon shielding during the stress-relief process.



Figure 3.8: The test artifacts (NIST parts) built with the five Powder Systems listed in Table 3.1.
(a): PS #1 (nominal).
(b): PS #2.
(c): PS #3.
(d): PS #4, the heavy oxidation is due to an accidental loss of argon shielding during the stress-relief.
(e): PS #5, the coloration for this In718 artifact is noticeably different than for the other, Ti64, artifacts.

It should be noted that the PS #5 build failed approximately 75% of the way through the

build height due to a machine error unrelated to either the part geometry or the powder being

used (see Section 5.2.6). This failure is visible as several partially delaminated layers in Figure

3.9. The PS #4 build experienced delamination and cracking on the NIST part (Figure 3.10), as

well as several of the cylinders. After analyzing data recorded by a custom autonomous powder

bed monitoring system (Chapter 4), it was determined that the cracking and delamination occurred sometime after the build completed, but before it was removed from the build chamber. It should be noted that while the cause of this cracking is unknown, similar cracking of samples built with PS #4 was also observed by Hengfeng Gu at NCSU during the course of the America Makes project.





Figure 3.9: A representative cylinder built with PS⁴³ #5. The EOS M290 spontaneously failed to raise the powder dispenser (Figure 1.4), resulting in several layers of insufficient powder spreading and eventual delamination.

Figure 3.10: Cracking (contiguous red oval) and delamination (dashed red oval) during the PS #4 build.

All of the small-scale positive features (Figure 3.11) on the NIST part were built successfully

in each Powder System tested, but there was a slight degradation in the resolution of the

smallest-scale rectangular (vertical) negative feature for PS #3, relative to the other Powder

Systems (Figure 3.12).

⁴³ The "PS 7" label on the cylinder image references the original Powder System numbering scheme used for the America Makes report [53]. Please refer to the following conversion key: thesis (America Makes), PS #1 (PS #8), PS #2 (PS #3), PS #3 (PS #4), PS #4 (PS #5), and PS #5 (PS #7).



Figure 3.11: An example of successfully produced vertical positive features, specifically those from the NIST part produced using PS #5.

Figure 3.12: Negative feature resolution on the NIST part built using PS #3; note the near merger of the two rectangles highlighted by the red circle.

3.3.2 Quantitative Analysis of the Powder Spreading Process

The images collected by the EOS M290's EOSTATE PowderBed module [105] during each build were analyzed using the ML/CV methodology presented in Chapter 4. Specifically, a Multi-scale Convolutional Neural Network (Section 4.4) was used to detect and classify anomalies located on the power bed throughout the height of the build. Figure 3.13 shows the occurrence of anomaly detections⁴⁴ as a function of build layer for all five Powder System builds. It is immediately evident that PS #2 and PS #3 have dramatically higher numbers of anomaly detections than the other three Powder Systems. It is also evident that a strong periodicity in the anomaly detections exists for PS #2 and PS #3 while a less dramatic periodicity exists for PS #4. Manual review of the powder bed images, as well as in-person observations during the build (Figure 3.7), suggest that these anomaly detections are triggered by clumps of dark particles that appear at the edges of the stripes (Figure 1.8) during fusion of the powder layer. These dark particles are not readily apparent in the builds using PS #1 or PS #5.

⁴⁴ Specifically, the anomaly detections reported in this section are a combination of detections of *super-elevation* and *part damage*, the descriptions for both of which can be found in Section 4.2.4.



Figure 3.13: The number of pixels (as a percentage of the cross-sectional part area) that the Multi-scale Convolutional Neural Network has identified as anomalous at each layer of the build.

A Fourier frequency analysis yielded a period of 5.37 layers/anomaly-peak for the PS #2 – PS #4 builds. The author hypothesizes that the periodicity is related to the default EOS M290 laser scan strategy (Figure 1.8) rotation of 67° every layer [41]. This rotation strategy results in a "near-full" rotation (335°) every five layers; observe that (5.37 layers/anomaly-peak) × (67°/layer) = 360°/anomaly-peak, i.e. a full rotation. A review of the EOSPRINT [45] build file (Figure 3.14) indicates that the anomaly detections peak when the edges of the stripes exposed during the previous layer are perpendicular to the *x*-axis. While it is possible that there is an increased degree of interaction between the recoater blade (which travels along the *x*-axis), the new powder spread, and the dark particles, it is likely that the periodicity is an artifact of lighting conditions within the build chamber. Specifically, the lighting source is placed such that it illuminates the powder bed from the right to the left along the *x*-axis (as in Figures 1.4, 1.5, and 3.14). As a result, the lighting contrast that highlights elevation variations on the powder bed surface will be greatest for features which are oriented perpendicularly to the recoater

blade direction, i.e. the anomalies are easier or harder to detect depending on their orientation relative to the illumination source.



Figure 3.14: First CAD geometry slice of the test artifacts as viewed in EOSPRINT. Note the vertical stripes (perpendicular to the illumination and recoater blade travel directions). Given the 67° rotation, the stripes will align with the *y*-axis every 5.3 layers.

Based on the relationship between the dark particles and the stripe edges, it is likely that a significant percentage of the dark particles are a combination of ejecta from the melt pool and powder particles entrained by the intense fluid flows within the vapor plume, commonly referred to as spatter [39], [106], [107]. Recent studies have suggested that spatter lying on top of the powder bed has the potential to negatively influence the fusing of subsequent powder layers and overall part quality [106]. Part quality for each Powder System is evaluated in the following two subsections.

3.3.3 Bulk Porosity of As-Built Cylindrical Samples

Figure 3.15 shows the pores detected and categorized in an example micrograph. Figure 3.16 shows the percentage of the of the bulk region that contains porosity of each type. Each reported Powder System measurement represents the average of the values determined for five cylinder samples. Average bulk porosity was measured to be less than 0.1% for all Powder

Systems. Cylinders built with PS #3 have the highest level of detected porosity (p-value⁴⁵ of 0.065), while the remaining Powder Systems (#2, #4, and #5) have porosity levels observably similar to that of the standard EOS powder (PS #1). The primary contributors to the increased porosity in the PS #3 samples are large, irregularly-shaped pores classified as the result of the lack-of-fusion mechanism. Given that the processing parameters were identical for the PS #1 – #4 builds (Table 3.3), the author suspects that some of this porosity may be the result of an interaction between the melt pool and the spatter particles accumulating at the edges of the stripes (Sections 3.3.1 and 3.3.2).



Figure 3.15: A visual representation of the pore-classifications for the PS #2 micrograph shown in Figure 3.5. Note that the effective circular diameters of the detected pores have been increased by a factor of eight to improve visibility.

⁴⁵ In a right-tail test the alternative hypothesis is that the mean of one population (PS #3) is larger than the mean of another population (in this case, PS #2 as it has the next highest level of bulk porosity).



Figure 3.16: The bulk porosity of the sectioned cylinders. The error bars represent 95% confidence intervals about the mean total porosity (the combined contribution of all four porosity types).

Also of interest, PS #4 (which experienced cracking, see Figure 3.10) and PS #5 (the only In718 Powder System) appear to have a higher percentage of porosity caused by small pores, which are considered to originate from the powder itself, as compared to the other three Powder Systems. Internal work by Ross Cunningham at CMU utilized a synchrotron to perform high resolution X-Ray Computed Tomography (μ CT) scans of both the cylindrical samples and a powder sample of each Powder System. The μ CT results also indicated higher percentages of porosity due to sub 20 μ m pores in in both the cylindrical and powder samples of PS #4 and, to an even greater extent, PS #5 (relative to the other three Powder Systems). For both PS #4 and #5, the powder samples contained on the order of 10 times more sub 20 μ m pores than the corresponding cylindrical samples.

3.3.4 Edge Roughness of As-Built Cylindrical Samples

Figure 3.17 shows an analysis of the 2D edge roughness of the cylinders. Each reported Powder System measurement represents the average of the values determined for the two build direction (*z*-axis) edges of each of five cylinders (i.e. ten total edges). Each measured edge was between 1 cm and 2 cm in length. While the roughness values for all of the Powder Systems are comparable, the use of PS #3 resulted in rougher test artifacts than the other Powder Systems (p-values⁴⁶ of 0.0023, 0.0017, 0.0052, and 0.023 for the *Ra*, *Rq*, *Rz*, and *Rsk* metrics respectively). Notably, the mean *Rsk* value was negative for all of the Powder Systems. Internal modeling work by Christopher Kantzos of CMU has demonstrated that surface valleys are a significant source of stress concentrations (potential crack initiation sites) under certain tensile loading conditions. It can therefore be inferred that a highly negative *Rsk* value may indicate that a part has a decreased resistance to fatigue failure.

⁴⁶ In a right-tail test the alternative hypothesis is that the mean of one population (PS #3) is larger than the mean of another population (PS #2 as it has the next largest roughness values). Note that a left-tail test was used for the *Rsk* metric.



Figure 3.17: The edge roughness of the sectioned cylinders. The error bars represent 95% confidence intervals about the mean.

Across the five Powder Systems, the *Ra* metric ranges from 16 μ m to 29 μ m, the *Rq* metric ranges from 21 μ m to 38 μ m, the *Rz* metric ranges from 94 μ m to 180 μ m, and the *Rsk* metric ranges from -0.13 μ m to -0.033 μ m. The roughness of cylinders built with PS #5 is statistically indistinguishable from the roughness of cylinders built with PS #1, the standard EOS powder (pvalues⁴⁷ of 0.16, 0.87, and 0.35) for the *Rq*, *Rz*, and *Rsk* metrics respectively). However the pvalue for the *Ra* metric is 0.062, indicating that PS #1 and PS #5 may be statistically distinguishable from each other according to this metric. Overall, the similarity between the PS #1 and PS #5 results suggest that in some cases a reasonably wide range of powder sizes (mean particle diameters ranging from 34 μ m to 58 μ m) could be used to fabricate L-PBF parts without an appreciable increase to the as-built edge roughness of the parts.

⁴⁷ In a two-tail test the alternative hypothesis is that the means of the two populations are not equal.

3.3.5 Correlations between Build Quality and Powder Particle Size

A general trend of worsening build and part quality as powder particle size increases is apparent throughout the results in Sections 3.3.2 - 0. For example, the largest powder (PS #3) consistently ranks the worst across all of the metrics examined. In Figure 3.13 the Multi-scale Convolutional Neural Network determined that the PS #3 build experienced the greatest occurrence of anomalies in the powder bed and both the bulk porosity and the edge roughness of the PS #3 cylinders are statistically higher than that of the other Powder Systems. The correlation between the build and part quality metrics with respect to mean particle size and maximum particle size (as defined by the upper 10% threshold, see Table 3.2) was quantified with a Pearson rank correlation test [108]. The results of the correlation test⁴⁸, specifically the correlation coefficients (ρ) and the associated p-values, are presented in Tables 3.5 – 3.7.

Table 3.5 demonstrates a correlation between the powder particle sizes and the powder bed anomalies detected by the Multi-scale Convolutional Neural Network. This correlation suggests that the Powder Systems with larger powder particles result in the generation of more (or more visible) spatter particles during fusion of the powder bed. Table 3.6 demonstrates a correlation between the powder particle size and several of the bulk porosity types. Note the lack of correlation in the case of the smaller porosity (attributed to the powder) and the case of the porosity attributed to the keyholing mechanism. This lack of correlation is expected as

 $^{^{48}}$ A Pearson rank correlation test quantifies how well the relationship between two variables follows a monotonic function. In this application of the correlation test, high correlation coefficients and low p-values indicate strong correlations between the particle size characteristics and the quality metrics. The correlation coefficient ranges from \mathbb{R} [-1, 1], with a value of 1 indicating a monotonically increasing relationship and a value of -1 indicating a monotonically decreasing (i.e. inverse) relationship [108]. Note that the reported correlation metrics are for the mean measurement values (where applicable).

native powder porosity is primarily related to the powder production method and keyholing porosity is primarily related to the L-PBF processing parameters (see Chapters 2 and 6). Finally, Table 3.7 demonstrates a correlation between the powder particle size and several of the roughness measures. Note that the greatest correlation is found for the *Rz* metric, possibly because the magnitude of this value may be driven by unfused or partially-fused powder particles adhered to the surface of the cylindrical samples. Also note that no correlation is found for the *Rsk* metric. As the *Rsk* value describes the presence of sharp surface valleys, it is not surprising that it is not significantly influenced by the powder particle size. Advantageously, this lack of correlation suggests that the use of Powder Systems with larger powder particles does not necessarily result in as-built surfaces that are more susceptible to crack initiation under fatigue loading conditions (see Section 3.3.4).

Interestingly, the correlations are markedly stronger with respect to mean particle size than with respect to maximum particle size. A potential contributing factor to this effect may be the use of effective powder layers (approximately 100 μ m thick) that are thinner than the largest powder particles (approximately 120 μ m in diameter). The powder layer thicknesses were determined based on the PSDs reported by the manufactures (Table 3.1) as powder characterization (Table 3.2) was not performed prior to the builds.

Table 3.5: Results from an application of the Pearson rank correlation test to the quantitative analysis of the powder spreading process. Table cells are shaded progressively darker for p-value cutoffs of 0.20, 0.10, and 0.05.

	Average Powder Bed Anomaly Percentage	Peak Powder Bed Anomaly Percentage	
	Particle Size (Mean)		
ρ	0.88	0.93	
p-value	0.049	0.020	
	Particle Size (Upper 10% Threshold)		
ρ	0.68	0.72	
p-value	0.21	0.17	

Table 3.6: Results from an application of the Pearson rank correlation test to the measured bulk porosity types. Table cells are shaded progressively darker for p-value cutoffs of 0.20, 0.10, and 0.05.

	Native Porosity	Keyholing Porosity	Lack-of-Fusion Porosity	Unclassified Porosity	Total Porosity
	Particle Size (Mean)				
ρ	-0.12	0.65	0.74	0.84	0.88
p-value	0.85	0.23	0.16	0.048	0.068
	Particle Size (Upper 10% Threshold)				
ρ	0.26	0.36	0.52	0.98	0.81
p-value	0.67	0.55	0.36	0.0043	0.099

Table 3.7: Results from an application of the Pearson rank correlation test to the measured edge roughness metrics. Table cells are shaded progressively darker for p-value cutoffs of 0.20, 0.10, and 0.05.

_	Ra	Rq	Rz	Rsk
	Particle Size (Mean)			
ρ	0.85	0.86	0.90	-0.50
p-value	0.068	0.061	0.038	0.39
	Particle Size (Upper 10% Threshold)			
ρ	0.59	0.61	0.69	-0.55
p-value	0.30	0.27	0.20	0.34

3.4 Discussion and Summary

In this chapter, four non-standard Powder Systems were successfully used to build a standard set of test artifacts in an L-PBF process. The non-standard Powder Systems spanned two alloys (Ti64 and In718), three powder production techniques (Gas Atomization, PREP, and HDH + P-S), and had PSDs with maximum particle sizes up to 2.4 times larger than the largest particles found in the EOS nominal Ti64 Powder System. As a control, the test artifacts were also built using the EOS nominal Ti64 Powder System.

The powder spreading and build qualities were evaluated qualitatively. Observations were made with respect to the resolution of a selection of the small-features on the test artifacts built with each Powder System. Post-build cracking of the PS #4 artifacts was observed, the cause of which is unknown at this time. Quantitative evaluation of the powder spreading process was performed using the powder bed monitoring methodology (Multi-scale Convolution Neural Network) presented in Chapter 4. The algorithm detections indicate that powder bed anomalies, in the form of dark particles, were more common in PS #2 – PS #4 than the nominal Powder System. Additionally, it was shown that the spatial distribution of the anomalies is related to the laser scan strategy, suggesting that the dark particles may be partially composed of spatter particles – further motivating the need for robust in-situ process monitoring methodologies such as those presented in Chapters 4 and 7.

As-built part quality was evaluated for each Powder System via measurements of bulk porosity and 2D edge roughness. The average bulk porosity remained below 0.1% for all of the Powder Systems, although the porosity is noticeably higher for PS #3. Based on the morphology

80

(large and irregular) of many of the pores in PS #3, it is conceivable that the increased porosity is related an interaction between the melt pool and the accumulated spatter particles on top of the powder bed. The edge roughness values of the cylindrical samples built with the nonstandard Powder Systems are similar to those of the cylindrical samples built with the nominal Powder System; note, however, that the edge roughness is statistically highest for the PS #3 sample relative to the other Powder Systems.

A rank correlation analysis revealed a trend of worsening powder layer and as-built part quality with increasing mean powder particle size. Specifically, correlations were found between mean powder particle size and the detected powder bed anomalies, some of the bulk porosity types, and some the of edge roughness measures. Advantageously, no correlation was found between powder particle size and the edge roughness measure most closely linked to the formation of near-surface stress concentrations. Interestingly, only weaker correlations were found between the various quality metrics and the maximum powder particle sizes. Despite the observed trend, the overall magnitudes of the quality metrics suggest that parts of acceptable quality (for many applications) can be built in L-PBF processes using non-standard Powder Systems with dramatically larger particle sizes than the Powder Systems currently recommended by the machine manufacturers.

- 4 Topic 3: Autonomous Powder Bed Anomaly Detection and Classification in an L-PBF Process using Machine Learning Techniques
- 4.1 Background and Literature Review

The applications best suited for Additive Manufacturing require a degree of part quality assurance and process reliability that are difficult to achieve with the systems currently on the market [2]. It is commonly recognized that implementation of in-situ process monitoring and closed-loop control is necessary to meet the stringent requirements of these applications [2]. In-situ process monitoring of builds has become a major research focus for the AM community over the last several years. Monitoring efforts for the PBF and DED AM processes have variously focused on detecting macro-scale flaws (e.g. part delamination and residual stress-induced warping) [109], [110], detecting micro-scale flaws (e.g. porosity), measuring temperature fields and histories [109], [110], measuring shielding gas quality [111], and understanding melt pool dynamics [109], [110]. An impressive range of sensor modalities have been explored including those enumerated in the remainder of this paragraph. High speed pyrometers and high speed thermal imaging to measure melt pool temperatures [112]–[114]. Low speed pyrometers and low speed thermal imaging to measure powder bed temperatures [34], [115]–[117]. Embedded thermocouples to measure build substrate temperatures [118]. High speed visible-light imaging (see Chapter 7 for a more detailed treatment) [81], [119]–[124], high speed X-Ray imaging [85], [125], and interferometric coherence imaging [126] to monitor melt pool, spatter, and vapor plume sizes and shapes. Strain gages to directly measure part distortion [118], [127]. And perhaps most recently, active [128], [129], passive [130], [131], and spatially resolved acoustic [132], [133] sensing to detect a variety of flaw signatures.

Many of the flaws in a final part, as well as the overall reliability and stability of the L-PBF build process, are directly related to interactions between the recoater blade and the powder bed. The various powder bed anomalies identified in the literature and studied in this chapter are summarized in Table 4.1. These anomalies range in severity from *recoater hopping* which may only indicate the onset of a more severe problem, to *super-elevation* which can impact the stability of the entire build. Some anomalies (such as *part damage*) may indicate flaws in the final part, while others, such as *recoater streaking*, suggest damage to the machine itself; further descriptions of the anomalies are provided in Section 4.2.4.

As a result, several groups have begun paying special attention to this stage of the L-PBF process using low speed visible-light imaging of the powder bed [134]–[140], in some cases in conjunction with flash-bulb illumination [141], [142], or structured light (i.e. fringe projection) [143]. Detection of *recoater streaking* has been explored by Craeghs et al. [136] and various methods for detecting *super-elevation* (albeit at a different size scale) have been proposed by Jacobsmühlen et al. [140]. Work by Foster et al. [141], Abdelrahman et al. [142], and Aminzadeh [138] demonstrates layer-wise geometric measurements and detection of general flaws via comparison of post-fusion visible-light powder bed images with the CAD model; similar work has also been performed by Cooke et al. [144] but for a Polymer Binder Jetting system. Finally, it is worth noting that L-PBF machine manufacturers (including EOS GmbH [145] and ConceptLaser GmbH [146]) are now releasing process monitoring solutions that include analyses of the powder bed. Unfortunately, such details as would be necessary to independently validate and improve upon the methodologies used by these systems are considered proprietary and are currently unavailable to the author.

83

Table 4.1: Descriptions of powder bed anomaly classifications and their respective color codes used throughout this chapter and Chapter 5. Additional descriptions of the anomaly types can be found in Section 4.2.4. Referring to the L-PBF background information provided in Section 1.1 may be useful.

Anomaly	Description	Overlay Color Codes	Non-Overlay Color Codes
Okay	No significant anomalies in the powder bed.	Transparent	Green
Recoater Hopping	Caused by the recoater blade striking a part, characterized by repeated vertical (parallel to the y-axis) lines [141].	Light Blue (Teal)	Light Blue (Teal)
Recoater Streaking	Caused either by the recoater blade dragging a contaminant across the powder bed or by damage to the recoater blade. Characterized by horizontal (parallel to the <i>x</i> -axis) lines [136].	Dark Blue	Dark Blue
Debris	Debris or other small to mid-sized discrepancies located in the powder bed but not directly over any parts [141].	White	Black
Super-Elevation	Occurs when a part warps or curls upwards out of the powder layer [140]. Typically the result of a buildup of residual thermal stresses [118] or swelling [147].	Red	Red
Part Damage	General classification for any significant damage to a part. Characterized by a variety of signatures.	Purple (Magenta)	Purple (Magenta)
Incomplete Spreading	Occurs when an insufficient amount of powder is repeatedly fetched from the powder dispenser [146]. Results in a lack of powder, the severity of which is highest nearest the powder collector.	Yellow	Yellow

The powder bed monitoring approaches reported in the literature have a number of critical limitations: In all of the reported work, either only one anomaly type is detected by a given algorithm, or all of the detected anomalies are treated as generic "flaws." Such an approach is inherently limiting with respect to future feedback control applications as different anomalies require different mitigation strategies (Section 5.3). Additionally, much of the reported work requires relatively high effective camera resolutions⁴⁹ and therefore cannot be easily used to monitor the entire powder bed [138]–[140] and only a few authors have demonstrated the ability to reconstruct layer-wise data for an entire build [141], [142]. Furthermore, all of the reported work has required the installation of additional hardware beyond what is provided by the machine manufacturers. Of greatest importance, however, is the reliance of the existing

⁴⁹ In this context, the "effective camera resolution" refers to the number of camera pixels covering a given area. It can be considered the inverse of the Field-of-View of each pixel.

work on human-created detectors⁵⁰. For example, line profiles [136] and segmentation [138], [141], [142] have been used extensively. Such a human-dependent approach is severely limiting, as time-consuming human intervention would be required to modify the anomaly detection algorithm(s) as L-PBF machines evolve, new material systems become available, and new part geometries interact with the powder bed in unique ways.

Fortunately, Machine Learning enables an algorithm to create its own detectors and models – allowing its anomaly detection capabilities to improve and evolve as new training data are provided. Indeed, it is for this reason that ML has become extremely popular in the Computer Vision community, though many of the methods are typically applied to entire images [148]. This presents a challenge as each powder bed image may contain hundreds of uniquely-identifiable anomalies. The problem of classifying multiple objects within a single image has been addressed by many groups who used methods including: Bag of Words applied to patch-wise classification [149], pixel-wise classification using a sliding window approach [150], and Multi-scale Convolutional Neural Networks [151]. It should also be noted that recent work by Jacobsmühlen et al. [140] used various ML techniques to detect *super-elevation* in post-fusion images with a very high effective resolution visible-light camera. Specifically, HOG/DAISY⁵¹ features were used to describe the imaged part geometries and models for predicting *super-elevation* were learned using SVM (Support Vector Machines), RF (Random Forests), and SGD (Stochastic Gradient Descent) techniques. Note that HOG/SIFT is conceptually

⁵⁰ In this context, "detectors" refers to algorithmic solutions to identifying and classifying powder bed anomalies, not hardware-based detectors (i.e. sensors).

⁵¹ In this usage, the DAISY descriptor was modified to be rotationally-invariant.

similar to DAISY and is discussed in Section 7.3; SGD is discussed in Section 4.4.6; SVMs are discussed in Section 7.3.6.

In this chapter, the six powder bed anomalies presented in Table 4.1 are automatically detected and classified in a commercially-available L-PBF process via analysis of post-spreading (as opposed to post-fusion) powder bed images. Two different supervised ML methods are used to perform this task. The first ML method utilizes a Bag of Words (or Keypoints) [44] approach similar in concept to the patch-wise object classification reported by Winn et al. [149]. The second ML method utilizes transfer learning [152] to retrain the existing AlexNet Convolutional Neural Network (CNN) [153] to classify the powder bed anomalies. In addition to retraining, the AlexNet CNN was also configured such that the input data structure is composed of multi-scale patches, that is, regions of the powder bed at varying size scales. Intriguingly, the use of CNNs with multi-scale patches has, to the author's knowledge, only been reported once in the literature – work by Shen et al. [151] leveraged this approach (which they referred to as MCNNs) to classify cancerous nodules in the human lung based on thoracic Computed Tomography (CT) scans.

Theoretical concepts relevant to the two ML techniques are covered and the specific methodologies used in both cases are presented in detail. The performances of each ML methodology are evaluated and compared. Specifically, the guessing accuracies are measured using both validation and testing data sets, the sensitivity to training database size is reported for the CNN, and approximate computational burdens are reported. Based on the comparison of the aforementioned performance metrics, and additional qualitative considerations, a

86

decision is made regarding the relative efficacy of each ML approach as it relates to the powder bed monitoring application.

A single build is then used as a case study to demonstrate the capabilities of the final ML algorithm. Through this case study, it is shown that the final ML algorithm provides valuable information regarding part quality and overall build stability. The work presented in this chapter is further extended in Chapter 5 in which numerous builds are analyzed using the final ML methodology and additional conclusions are drawn regarding the influence of part geometry and process parameters on build stability and powder layer deposition. The work presented in this chapter in this chapter was not directly supported by any funding agencies in the public, commercial, or not-for-profit sectors.

4.2 Experimental Design and Methods

4.2.1 Programming Environment

Unless otherwise noted, all software was developed within the MATLAB R2015a, R2016a, or R2017a environments. The above MATLAB versions also included the following add-on packages: the Image Processing Toolbox, MATLAB Compiler, the Neural Network Toolbox, and the Statistics Toolbox.

4.2.2 Powder Bed Camera

All of the work presented in this chapter is performed on an EOS M290 L-PBF machine (Figure 1.4) and uses only the stock powder bed camera and lighting configuration. Images of the build plate and powder bed are taken through a viewport located (almost) directly above

87

the build chamber. Grayscale images with a resolution of 1280 pixels \times 1024 pixels are automatically captured immediately after a new powder layer is spread; the images are later accessible through EOSTATE [45]. During printing, the powder bed is illuminated by a single bank of white LEDs on the right-hand side of the build chamber. Figure 4.1 shows an example raw image taken by the powder bed camera.

4.2.3 Image Pre-Processing

The raw images (Figure 4.1) captured by the EOS M290 present several difficulties that prevent their direct usage in a ML-based algorithm. Fortunately, the camera mounting and lighting conditions remain consistent throughout a build as well as between different builds, so many of the required image adjustments can be greatly simplified.

Out of the necessity of avoiding the laser optic train, the powder bed camera is mounted such that its axis is not parallel to the normal vector of the build plate (*z*-axis). The resulting distortion is corrected using a fully-constrained Homography matrix [154] to apply an affine warp to the raw image such that a square build plate in the initial image will appear square in the final image. Because the camera positioning and orientation are fixed, measurements of a powder-free build plate (within the camera's field of view) were performed manually and no automatic fiducial (e.g. corner) detection is implemented to inform the Homography matrix. For the case of a non-fixed powder bed camera, work by Jacobsmühlen et al. proposes the use of robust fiducials exposed on the powder bed itself [134]. The image is then cropped to include only the region of the powder bed directly above the build plate. The spatial resolution (not synonymous with resolving power [155]) of the camera setup is between⁵² 290 μ m/pixel and 340 μ m/pixel. After the described warping and cropping, the image is 900 pixels × 900 pixels with each pixel representing a 280 μ m × 280 μ m field of view.

The highly directional side lighting of the powder bed increases (compared to top lighting) the contrast of any 3D features (e.g. hills and valleys). Overall contrast within the image is further increased by a remapping of the pixel intensities such that 1% of the pixels at each intensity extreme are saturated [156]. A side effect of the lighting configuration is a haloing effect (visible in Figure 4.2) in the images that is detrimental to the ML training process. To remedy the uneven lighting conditions, an anomaly-free powder bed image is used to generate a baseline intensity mask. Stochastic noise present in the mask is reduced using a Gaussian filter with a standard deviation of 8 pixels. This mask (Figure 4.2) is applied to each future powder bed image to levelize the lighting across the powder bed. Figure 4.3 shows a fully pre-processed image which has been warped, undergone a contrast adjustment, and had the lighting adjustment mask) should be expected to vary across EOS M290 machines due to manufacturing and assembly differences.

⁵² The existence of a range of resolutions is the result of the misalignment between the camera axis and the normal vector of the build plate.



Figure 4.1: Raw powder bed image collected by the EOS M290.



Figure 4.2: The illumination mask generated during the calibration process. The center of the powder bed is brightest and requires the least illumination correction, while the left-upper and left-lower edges require the most significant illumination correction.



Figure 4.3: Figure 4.1 after pre-processing.

4.2.4 Powder Bed Anomaly Types

The powder bed anomaly classifications chosen for this work are briefly introduced in Table 4.1 and further described in this subsection. Importantly, the operations of the ML methodologies presented in Sections 4.3 and 4.4 are not dependent upon the specific anomaly classifications chosen by the author; the ML approach only requires that the chosen anomaly types are self-consistent and sufficiently distinct from each other. Of course, a sufficiently large number of training examples must be available for each anomaly type. The specific anomaly classes (and their nomenclature) used in this work were chosen based on existing studies in the literature, the experiences of the author with operating the EOS M290, and feedback from multiple internal and external users of the developed powder bed monitoring methodologies. For the purposes of the consolidation heuristics described in Section 4.3.6 and the discussion in Section 5.3, the powder bed anomaly types in order of increasing severity are: *recoater hopping, recoater streaking, debris, super-elevation, part damage,* and *incomplete spreading.*

Recoater hopping typically occurs when the recoater blade (relatively) lightly strikes a part just below the powder layer. Such a strike results in a periodic "chatter" of the recoater blade which is visible as repeated vertical lines in the powder bed as shown in Figure 4.4a. *Recoater streaking* occurs either when the recoater blade is damaged (i.e. "nicked") or when the recoater blade drags a piece of debris or a clump of powder across the powder bed. *Recoater streaking* is visible as individual horizontal lines in the powder bed (Figure 4.4b). In many cases, these horizontal lines are only a few pixels wide in the powder bed camera image – increasing the challenge of robust detection. The *debris* classification encompasses many of the general disturbances to the powder bed that are not located directly over a part. Figure 4.4c shows several examples of *debris*.

Super-elevation occurs when a part visibly warps up above the powder layer [140]; the warping is often the result of a buildup in residual thermal stresses⁵³ within the part [118] or swelling [147]. Super-elevation anomalies have a wide range of appearances due to the wide range of extant part geometries, but all generally contain long edges at varying orientations. Figure 4.4d shows several representative examples of super-elevation. Part damage encompasses significant disturbances to the powder bed directly over a part and generally indicates damage to a part caused by a major strike by the recoater blade. Unlike superelevation, long edges are typically not present as the part has sustained damage; Figure 4.4e shows several examples of part damage. Finally, incomplete spreading occurs when insufficient powder is fetched from the powder dispenser and spread across the build plate⁵⁴. This results in large disturbances to the powder bed, which initially occur near the powder collector (lefthand side of the powder bed, Figure 1.4). Over the duration of a build, disturbances due to incomplete spreading may encroach further and further into the powder bed; Figure 4.4f shows an example of *incomplete spreading*. Finally, any regions of the powder bed not containing one of the six anomalies described above is considered to be *okay*.

⁵³ The large temperature gradient between the liquid melt pool and the EOS M290 build chamber (typically held at between 35 °C and 200 °C [45]) results in rapid solidification of the molten material and commensurately rapid and "non-uniform thermal expansions and contractions" [184] which induce residual stress within the part. Over the course of a build, the magnitude of the residual stress may exceed the local yield stress of the part, resulting in deformation – often expressed by the part "curling" or "warping" up above the powder bed [184].

⁵⁴ The amount of powder spread each layer is typically referred to as the "dosing factor."



Figure 4.4: Representative examples of the six different powder bed anomaly classes chosen by the author. Note that the relative sizes between the anomalies have been preserved. Specifically, the anomalies are: (a) *recoater hopping*, (b) *recoater streaking*, (c) *debris*, (d), *super-elevation*, (e) *part damage*, and (f) *incomplete spreading*.

4.2.5 Extraction of Part Geometry Information

To improve the utility of autonomous detection and classification of powder bed images, information about the nominal Computer Aided Design (CAD) geometry of the parts being built, as well as their location and orientation on the build plate, are overlaid on top of the powder bed images. This information enables the 3D reconstruction of parts with any anomalous regions highlighted (Section 4.6.5). Furthermore, as demonstrated throughout Chapter 5, this information also allows for the study of the influence of part geometry on powder spreading and process stability. Finally, part location on the build plate is a factor in certain anomaly classifications (as discussed in Section 4.2.4) such as *debris*, *super-elevation*, and *part damage*.

All of the required part information is contained within the EOSPRINT slice (.SLI) files [45]. Because the data within the slice files are encrypted [157], a CV approach was used to extract information from within the EOSPRINT v1.3 [45] environment. First, a Microsoft[®] Windows[®] 10 script, written using Autolt [158], automatically takes a screenshot of each layer of the build as it is displayed in EOSPRINT; Figure 4.5 shows an example screenshot. After all of the screenshots are collected, the build area is identified based on the white grid markings (Figure 4.5), the extraneous information is removed via cropping, and binary segmentation of the parts is performed based on the standard color pallet used by EOSPRINT. At this stage, the binary image is rescaled with bicubic interpolation [159] such that it is also 900 pixels × 900 pixels in size and can be directly overlaid on top of the warped powder bed images (Section 4.2.3). Finally, three dilation operations followed by three erosion operations are applied to the binary image in order to, respectively, "fill in" and remove noise from the part segmentations. Both the dilation and erosion operations utilize a 3 pixels × 3 pixels window. Figure 4.6 shows a final binary image encoding the part locations on the build plate.



Figure 4.5: A screenshot of a single layer as displayed in the EOSPRINT environment. The build area is indicated by the white grid and the part locations are shown in green and blue.



Figure 4.6: The final binary image encoding the part geometries and their locations on the build plate.

The layer-wise part location information for an entire build is converted from a 3D point cloud to a compressed data format in which a number triplet encodes the location of each "on-part" pixel in 3D space while the locations of the "off-part" pixels are not explicitly stored.

Because the part location information is always used in a layer-wise fashion, the speed of data decompression is improved by independently storing the location of the first on-part pixel in each layer within the compressed data structure. Explicit registration (alignment) between the powder bed images and the part CAD geometries, as discussed by Abdelrahman et al. [142], is not performed. Interestingly, another CV approach to slice extraction for AM has been presented by Vaidya et al. [160], although it was used for the purpose of converting NURBS-defined CAD geometry directly into slices without the use of an intermediate STL (Stereolithography) file format.

4.3 Bag of Words (BoW) Methodology and Theory

4.3.1 Overview

The methodology presented in this section is an application of a widely-used ML technique, known as Bag of Words (or Keypoints) (BoW) [44], often applied to CV problems. In this implementation, the training data consist of relatively small *patches* or regions of powder bed images which have been labeled by a human based on the anomalies they contain. While the BoW technique can be applied to multiple *feature* types, the author chose to use *filter responses* for their ability to preserve scale information (i.e. the size of a potential anomaly influences its *filter response*). This behavior is in contrast to the SIFT *features* described in Section 7.3. This section is intended to provide an overview of this methodology along with relevant ML and CV theory. Figure 4.7 is a flowchart of this ML methodology and is referred to extensively throughout this section.



Figure 4.7: Flowchart of the implementation of the BoW ML technique discussed in this section.

4.3.2 Selection of the Training Data

Each powder bed image may contain hundreds of distinct examples of different anomalies and *okay* regions. For this reason, training of the BoW ML algorithm is performed using image *patches*: sub-regions of the full powder bed images similar in concept to those shown in Figure 4.4. To develop the training database, a human manually selects rectangular image *patches* from multiple powder bed images captured during multiple builds. Note that, in general, only a subset of all of the possible patches within a powder bed image are selected for training. Also note that while the exact size of the training image *patches* is not constrained, they were selected such that they are similar in size to the *patches* used by the algorithm for anomaly classification (see Section 4.3.5). Prior to *patch* selection, the powder bed images are preprocessed as described in Section 4.2.3. A human labels each *patch* with a "ground-truth" anomaly classification (Section 4.2.4). The *patches* and their attached labels are stored in a database for access by the BoW ML algorithm. The training database includes a total of 2,402 image *patches*, composed of 1,040 *okay patches*, 264 *recoater hopping patches*, 228 *recoater streaking patches*, 187 *debris patches*, 314 *super-elevation patches*, 264 *part damage patches*, and 105 *incomplete spreading patches*.

4.3.3 Filter Bank

This implementation of the BoW ML algorithm extracts *features* using a set of *filters*. In this context, a *filter* refers to a discretized, 2D function or pattern; three examples of *filters* can be seen within the representative *filter bank* shown in Figure 4.7a. These *filters* are "passed over," or convolved with, an image (either a training image *patch* or a powder bed image). A convolution operation, shown graphically in Figure 4.8, is simply the summation of the element-wise multiplication of two matrices. The output of the convolution operation is an image with the same dimensions as the original image. The value of each pixel of the output image is the *response* of the original image to the convolution with the applied *filter* – its value depends upon the values of the corresponding pixel and its surrounding pixels in the original image.



Figure 4.8: A graphical representation of 2D convolution. In this example, a vertical edge *filter* is convolved with two regions of an image. One region (solid magenta box) contains a vertical edge and produces a high *response* while the other region (dashed magenta box) does not include a vertical edge and produces a low *response*.

Different *filters* produce either stronger or weaker *responses* depending on the distribution of the information encoded in the original image. For example, a Gaussian *filter* (e.g. filter 1 in Figure 4.7a) produces a strong *response* when convolved with pixels in the vicinity of a dark blob, while an asymmetric first derivative *filter* (e.g. filter 2 in Figure 4.7a) produces a strong *response* centered on any dark, vertical edges. The *response* of each pixel to each *filter* in the *filter bank* is stored in a vector, represented by the vertical bars in Figure 4.7b; the three colored segments of each bar represent the *responses* to the three *filters* in the representative *filter bank*. These *response* vectors are the format for the *features* on which this ML algorithm operates.

The *filter bank* used in this section contains a total of thirty-seven *filters*, each of size 25 pixels \times 25 pixels, and enumerated in Table 4.2. Multiple trials were performed to determine an effective *filter bank* size and composition. To mitigate the impact of edge effects on the training process, the borders of the image *patches* are padded with the pixels surrounding that *patch* while the borders of the powder bed images are padded with replicated pixels [161].

Table 4.2: A brief description of the composition of the *filter bank*.

Filter Type	Description	Num. of Size Scales/Variants
Gaussian	Standard <i>filters</i> , designed to detect blobs.	3
Uniform Averaging Disk	Standard <i>filters</i> , designed to detect blobs. Their responses are generally stronger than those of a Gaussian <i>filter</i> for blobs of particularly uniform intensity (i.e. darkness).	1
Difference of Gaussian (DoG)	Standard <i>filter</i> , designed to detect edges at all orientations.	3
Oriented Edge Detectors	Standard, asymmetric first derivative <i>filters</i> designed to detect edges at specific orientations (0° and 90°).	6
Oriented Line Detectors	Non-standard combinations of oriented edge detectors designed to detect lines at specific orientations (26°, 45°, 64°, 116°, 135°, and 154°). Particularly effective at detecting <i>super-elevation</i> anomalies.	18
Streak Detectors	Non-standard combinations of oriented edge detectors designed to detect <i>recoater streaking</i> and <i>recoater hopping</i> anomalies.	4
Gabor [162]	Standard <i>filters</i> designed to produce strong responses when convolved with an image containing spatially-encoded frequency information ⁵⁵ .	2

4.3.4 Training

The *filter bank* described in Section 4.3.3 is passed over every training image *patch*, such that there is a *response* vector for every pixel in every image *patch*. No subsampling of the training data is performed, i.e. all of the *response* vectors are included in the training process. *Response* vectors with similar values in each element are grouped together using a standard k-means unsupervised clustering algorithm [163], represented by Figure 4.7c. For this work, cluster initialization was performed using random seeding, with preference given to a uniform spacing between clusters. The requested number of clusters was systematically increased until the final anomaly classification results ceased to noticeably improve; the final clustering produces 100 groups. Cluster seeding is repeated 100 times to reduce the chance of the algorithm converging to a poor solution; e.g. a shallow local minimum instead of a global, or at least a deeper local, minimum.

⁵⁵ Only the numerically real component is considered in this implementation.

Each group is represented by a mean *response* vector. The 100 mean *response* vectors are commonly referred to as *visual words*, and are stored in a *dictionary*, represented by Figure 4.7d. The *visual words* are the *features* that will be searched for in future data sets (i.e. new powder bed images). After the *dictionary* is constructed, the *filter bank* is again convolved with each training image *patch*. But this time the *filter response* vectors at each pixel are matched to the closest (pair-wise distance [164]) *visual word* in the *dictionary* (Figure 4.7e). For each training image *patch*, the percentage of pixels matched to each *visual word* is calculated. This information can be represented by a histogram (Figure 4.7f). As it semi-uniquely identifies each image *patch*, it is often referred to as a *fingerprint*. The *fingerprint* of each training image *patch* is stored in a table (Figure 4.7g). Ideally, training images containing similar anomalies will have dissimilar *fingerprints*, while training images containing different anomalies will have dissimilar *fingerprints*. The final output of the training process is a table containing 2,402 *fingerprints* (one for each training image *patch*) that are each 100 elements long. The corresponding ground-truth anomaly labels for each *fingerprint* are stored in the training database (Section 4.3.2).

4.3.5 Patch-wise Classification

The steps described in Section 4.3.4 are only performed during training. During the classification of new data, the *filter bank* is convolved with the entire powder bed image and each pixel is assigned its closest-matching *visual word*. The layer image is then broken up into *patches* (Figure 4.7h). Performing the convolution operations on the entire powder bed image prior to *patch* delineation is computationally advantageous as it avoids duplicate computations that would otherwise be performed on the *patch* padding (Section 4.3.3).

Unlike the training *patches* in Section 4.3.2, the sizes of these anomaly classification *patches* are strictly defined. The *patches* are rectangular but vary in size and aspect ratio to better detect specific anomalies. A total of three different *patch* types are used: a 20 pixels \times 20 pixels square that is expected to detect most anomaly types, a 10 pixels \times 40 pixels rectangle designed to detect *recoater streaking*, and a 100 pixels \times 100 pixels square designed to detect *incomplete spreading*. Note that these anomaly classification resolutions correspond to minimum classifiable powder bed anomaly sizes of 5.6 mm \times 5.6 mm, 2.8 mm \times 11 mm, and 28 mm \times 28 mm, respectively (refer to Section 4.2.3). A *fingerprint* for each *patch* is generated and compared (using a binary singleton expansion function [164]) to the table of *fingerprints* (Figure 4.7i). It is at this stage that the labels associated with each *fingerprint* in the table are recalled from the training database.

The three training *fingerprints* closest to the *fingerprint* of the *patch* are considered during this step of the classification process. Specifically, the top three matches are weighted according to their respective degree of agreement and the *patch* is classified as the anomaly with the highest total weighting among the top three matches. For example, if the top match is for *recoater hopping*, but the next two matches are for *recoater streaking*, the algorithm would classify the patch as *recoater streaking* if the second and third matches have a stronger combined agreement than the first match individually. This approach was experimentally shown to produce more accurate results, possibly by mitigating the impact of over-fitting [165]. Additionally, if any of the top three matches are an *okay* case, the *patch* is always classified as *okay*. This restriction reduces the number of false anomaly classifications (i.e. false positives), which were deemed more problematic than false negatives as machine operators may be

reticent to use an overly-sensitive powder bed monitoring algorithm. The final anomaly classification for that *patch* is then applied to every pixel within that *patch*.

4.3.6 Layer-wise Classification and Contextual Heuristics

The overlapping results from the three *patch* type analyses are combined with relevant part geometry information (Section 4.2.5) using a series of contextual heuristics to determine the appropriate anomaly classification for each pixel in the powder bed image (Figure 4.7j). The inclusion of the heuristics layer allows information about the location of the anomaly detections with respect to the build plate, the parts, and the surrounding anomaly detections to factor into the final anomaly classification decisions made by the algorithm. Each of the heuristic rules is discussed in the following two paragraphs.

If the initial classification of a pixel not lying on top of a part (as defined by the extracted part geometry information) is either *part damage* or *super-elevation*, the classification is converted to the *debris* category as it can be visually similar to the aforementioned anomalies, which, by definition (Section 4.2.4), can only occur on top of a part. Similarly, any pixels lying on top of a part and initially classified as *debris*, are converted to the *part damage* category. Because *incomplete spreading* anomalies are detected with the largest scale *patches*, not all of the pixels labeled as *incomplete spreading* may be truly anomalous. To increase the chances that as-built part quality is accurately reflected by the BoW algorithm output, any pixels lying on top of a part and initially labeled as *incomplete spreading* are retroactively un-labeled. In other words the results from the largest *patches* are not included in the final, multi-*patch* consolidation process for the pixels lying on top of parts. An additional constraint is imposed
upon the *incomplete spreading* category to reduce the occurrence of false positives: Extremely large disturbances to the powder bed may appear visually similar to *incomplete spreading* despite being the result of a different mechanism (such as a severe impact of a recoater blade with a part). Therefore, *incomplete spreading* detections are ignored if there are no *incomplete spreading* detections near the left-hand edge of the build plate (see Section 4.2.4 for reasoning).

At this point in the implementation of the heuristics, the results from the three different *patch* types are combined. Because the *patch* results overlap and may be in disagreement, consolidation is effected by classifying each pixel as the highest possible severity anomaly, according the severity ranking listed in Section 4.2.4. After consolidation, the false positive rates for *recoater hopping* are reduced by confirming that multiple detections (a minimum of 80 pixels) are present in a vertical line (i.e. parallel to the *y*-axis). A similar approach is pursued for recoater streaking, except multiple detections (a minimum of 80 pixels) in a horizontal line (i.e. parallel to the *x*-axis) are required. To reduce the computational burden of the heuristic layer, many of the context rules are implemented by convolution of a mask of the initial anomaly detections with various templates and filters. In this case, the filters are conceptually similar to, but not the same as, the *filters* described in Section 4.3.2.

Finally, the layer-wise anomaly classifications for an entire build are converted from a 3D point cloud to a compressed data format in which a number quadruplet encodes the type of anomaly and its location in 3D space while the locations of the *okay* pixels are not explicitly stored. Because the anomaly classification information is always used in a layer-wise fashion,

104

the speed of data decompression is improved by independently storing the location of the first anomalous pixel in each layer within the compressed data structure.

4.4 Multi-scale Convolutional Neural Network (MsCNN) Methodology and Theory

4.4.1 Overview

The methodology presented in this section makes use of a pre-trained Convolutional Neural Network (CNN). CNNs have become increasingly popular in the CV community as they typically demonstrate high classification accuracies while leveraging prior knowledge about the input data (i.e. that the data are images) in order to simplify the algorithm architecture and reduce the training burden [153]. CNNs are an example of Deep Learning and require large data sets (on the order of $10^5 - 10^6$ samples) and significant computational resources to train [152]; fortunately, transfer learning allows for a pre-designed and pre-trained CNN to be applied to a unique classification problem [152]. Specifically, the version of the AlexNet CNN architecture, originally developed by Krizhevsky et al. [153], implemented by MATLAB [166] was applied to the patch-wise classification of powder bed anomalies. The AlexNet architecture and transfer learning are discussed further in Sections 4.4.3 - 4.4.5 and 4.4.6, respectively. As will be discussed in Section 4.4.3, the input data for AlexNet are typically color images, while the data structures in this work are grayscale patches at multiple size scales. For this reason, the ML algorithm presented in this section is referred to as a Multi-scale CNN (MsCNN) - a nomenclature borrowed from Shen et al. [151]. Figure 4.9 is a flowchart of this ML methodology and is referred to extensively throughout this section.



Figure 4.9: Flowchart of the implementation of the MsCNN ML technique discussed in this section. For clarity, architecture relating to the implemented parallelization has been omitted, as has an additional rectifier layer between (f) and (g).

4.4.2 Selection of the Training Data

Each powder bed image may contain hundreds of distinct examples of different anomalies and *okay* regions. For this reason, training of the MsCNN ML algorithm is performed using image *patches*: sub-regions of the full powder bed images similar in concept to those shown in Figure 4.4. To develop the training database, a human manually selects square image *patches* from multiple powder bed images captured during multiple builds. Note that unlike the BoW approach (Section 4.3.2), the size of each image *patch* used by the MsCNN is pre-defined as 25 pixels × 25 pixels. Therefore, there are 1,296 image *patches* in each powder bed image. As will be addressed in the following subsection (4.4.3), the input data structure for the MsCNN consists of a *multi-scale patch* composed of three related *patches*, each at a different size scale. Specifically, the "Level 1" *patch* is the aforementioned 25 pixels × 25 pixels region. The "Level 2" *patch* is a 100 pixels × 100 pixels region with the same center as the Level 1 *patch*. Finally, the "Level 3" *patch* is the entire 900 pixels \times 900 pixels powder bed image resized to 200 pixels \times 200 pixels using bicubic interpolation [159]. An example of each *patch* level is shown in Figure 4.9. Note that the powder bed images are pre-processed as described in Section 4.2.3 prior to *multi-scale patch* selection.

The observant reader may notice that centering a Level 2 *patch* on a Level 1 patch at the edge of the powder bed image is non-trivial, as the centered Level 2 *patch* will extend beyond the boundary of the powder bed image. This issue is addressed by padding the entire powder bed image with symmetric pixels [167] out to a distance equal to half of the Level 2 patch size (i.e. 50 pixels) at each edge. Furthermore, as first noted in Section 4.3.3, the use of small *patches*, relative to the *filter* size (i.e. the region over which a *feature* is extracted) requires that edge-effects be carefully considered. To mitigate the influence of these edge-effects, each *patch* is padded by 13 pixels on a side. In regions away from the edges of the powder bed image, padding is composed of the pixels surrounding that *patch* while at the borders of the powder bed image *patches* are padded with symmetric pixels [167]. This padding occurs subsequent to the padding applied to the edge-case Level 2 *patches* and results in padded *patch* sizes of 51 pixels × 51 pixels, 126 pixels × 126 pixels, and 226 pixels × 226 pixels for the Level 1, Level 2, and Level 3 *patches*, respectively. The choice of the padding size (13 pixels) was informed by the *filter* size and *filter* stride⁵⁶ of the first convolutional layer (Section 4.4.4).

A human applies each "ground-truth" anomaly classification (Section 4.2.4) to each *multiscale patch*. Note that the ground-truth anomaly classification only indicates the anomaly

⁵⁶ The stride of a *filter* specifies the spatial distance between the centers of the convolution operations.

present within the region covered by the Level 1 *patch*. The *multi-scale patches* and their attached labels are stored in a database used by the MsCNN algorithm during training (Section 4.4.6). The training database includes a total of 10,071 *multi-scale patches*, composed of 3,827 *okay patches*, 1,896 *recoater hopping patches*, 527 *recoater streaking patches*, 666 *super-elevation patches*, 1,297 *disturbance*⁵⁷ *patches*, and 1,858 *incomplete spreading patches*. The training *multi-scale patches* were extracted from a total of 89 powder bed images captured during a total of 14 builds.

4.4.3 Input Layer

All CNNs operate on data stored in an input layer (Figure 4.9a), or more strictly, an input volume of size width \times height \times depth [168]. The input layer of the AlexNet CNN was originally designed to operate on color images from the ImageNet dataset [169] and is of size 227 pixels \times 227 pixels \times 3 pixels, where the depth spans the three color channels (red, green, and blue) [153], [166]. When applying transfer learning to a pre-trained CNN, the CNN architecture, including the size of the input layer, must remain unchanged.

Typically, grayscale images (such as those captured by the powder bed camera) would be inserted into the AlexNet input layer by duplication of the grayscale data across all three color channels. This implementation is not ideal as redundant calculations will be performed on the input layer throughout the depth⁵⁸ of the CNN. Therefore the author chose to utilize the additional color channels to encode multi-scale information in the form of the *multi-scale*

⁵⁷ The *disturbance* anomaly classification encompasses both the *debris* and *part damage* anomaly types, see Section 4.4.7 for more details.

⁵⁸ In this usage, "depth" refers to the set of layers composing the CNN [168].

patches introduced in the previous subsection. Note that each Level patch is resized to 227 pixels \times 227 pixels using bilinear⁵⁹ interpolation [159] before being inserted into one of the color channels. This approach was inspired by the utility of the contextual heuristics and variously-sized *patches* implemented in the BoW methodology (Sections 4.3.5 and 4.3.6). Broadly, the contextual heuristics and multiple *patch* sizes allowed information about the size of the powder bed disturbances, adjacent anomalies, and the overall state of the powder layer to inform the layer-wise anomaly classifications. The goal of the *multi-scale patches* structure is to allow the CNN to learn those contextual relationships as opposed to the methodology relying on human-designed heuristics. Refer to Section 4.5.3 for validation of this desired behavior.

Specifically, the Level 2 *patch* is designed to encode information about the size of the disturbance on the powder bed and the region surrounding the Level 1 *patch*. Its size was chosen based on the author's experience with the 100 pixels × 100 pixels *patch* used by the BoW methodology to classify *incomplete spreading*, as well as the 80 pixel minimum length requirements used to reduce false classifications of *recoater hopping* and *recoater streaking*. The Level 3 *patch* is designed to encode information about the overall powder layer. Specifically, it is intended to serve the same purpose as the heuristic employed to reduce false classifications of *incomplete spreading* based on the presence (or lack-there-of) of significant disturbances on the left-hand edge of the powder bed. Additional reasoning for the *multi-scale patch* structure and the motivating contextual heuristics can be found in Section 4.2.4.

⁵⁹ In this application, bilinear interpolation is only slightly less accurate than bicubic interpolation but allows for a significant reduction in computation time.

4.4.4 Hidden Layers

Once the data are stored in the input layer, mathematical operations are applied to the data in a sequence of "hidden layers," so named because the operations learned by the CNN during training (i.e. the *feature* extraction tools) [43] and their outputs are difficult for humans to interpret and often non-trivial to describe [170]. The AlexNet CNN, and therefore the MsCNN used in this work, has a total depth of 25 layers [153], [166], 21 of which are considered hidden for the purposes of this subsection. This subsection broadly describes the various types of hidden layers used by the MsCNN; for a full accounting of the AlexNet architecture please refer to [153], [166].

As shown in Figure 4.9, the data stored in the input layer are first operated on by a convolution (CONV) layer; or more strictly, the results (or *responses*) of the convolution operations are stored in the CONV layer (Figure 4.9b). The convolution operations extract *features* using *filters*, conceptually identical to those introduced in Section 4.3.3. Critically, these *filters* are not chosen by a human, rather they are learned by the CNN during training (Section 4.4.6). For this reason, one may consider the CONV layer to be an optimized *filter bank*. Interestingly, the *filters* learned by CNNs for the first CONV layer are typically highly similar regardless of the specific classification application [152], [168], indeed this commonality provides the basic justification for the use of transfer learning [152].

The *filters* used in the first CONV layer of the MsCNN are of size 11 pixels \times 11 pixels \times 3 pixels [153], [166]; as is typical, each *filter* operates on the full depth of the layer (i.e. all three channels of the input layer) [168]. The size of the *filter* specifies the area of the input data over

110

which the convolution is performed while the "stride" of the *filter* specifies the spatial distance between the centers of the convolutions. In other words, for a stride of one, the convolution area moves one pixel in a given direction between operations (the *filters* in Section 4.3.3 are implemented with a stride of one). In the first CONV layer of the MsCNN, the stride is four [153], [166], i.e. the convolution area moves four pixels in a given direction between operations. A larger stride reduces the dimensionality of the CONV layer (a boon to computation time) but reduces the spatial resolution at which *features* are extracted (potentially detrimental to classification accuracy). The convolution operations result in a data volume with a depth equal to the number of *filters* and a width (W_o) and height given by (4.1). The volume of the first CONV layer in the MsCNN is 55 pixels × 55 pixels × 96 pixels [153], [166]. Because these *filters* operate through the depth of the input data volume, they are often referred to as *kernels*⁶⁰ [168].

$$W_o = \left(\frac{W_i - F + 2P}{S}\right) + 1 \tag{4.1}$$

Where W_o is the output layer width (or height), W_i is the input layer width (or height), F is the spatial width (or height) of the *kernel*, S is the stride of the *kernel* in the width (or height) direction, and P is the number of padding pixels explicitly used during the convolutions of the input data. Note that in the first CONV layer of the MsCNN, P is set to zero [153], [166] and padding is handled a priori (Section 4.4.2). Also note that the output size of a *pooling* operation can also be computed using this equation.

⁶⁰ Some readers may be familiar with the "human brain" analogy often used to describe the operation of CNNs. In this analogy, each *kernel* may be considered a *neuron* viewing a given *receptive field* (area) of the input data volume. The output (*response*) of each *neuron* may be considered an *activation* which is stored within the volume of the CONV layer.

Substantial work by the ML community has found that CNN training performance can be greatly improved by applying a non-linear positive-feedback⁶¹ operation to the output of each *kernel* [153], [171]. While hyperbolic tangent functions are often applied to the *kernel* outputs, Krizhevsky et al. [153] and others have determined that far superior training speeds can be obtained through the use of Rectified Linear Units (ReLU) which are defined in (4.2). The first ReLU layer in the MsCNN is represented schematically in Figure 4.9c. Note that the ReLU layer does not alter the size of the data volume, i.e. the output of the first ReLU layer in the MsCNN is of size 55 pixels × 96 pixels.

$$ReLU(\phi) = \max(0, \phi) \tag{4.2}$$

Where *ReLU* is the output of the ReLU operation and ϕ is the output of the *kernel*, i.e. the *response* of the convolution.

After rectification, Krizhevsky et al. [153] apply a Local *Response* Normalization (LRN)⁶²; in the MATLAB implementation of AlexNet this normalization occurs channel-wise [166]. The purpose of LRN is to further accentuate the spatially-local maximum *kernel responses*, thereby increasing the detection sensitivity of the CNN for spatially-small *features* [172]. Krizhevsky et al. [153] found that the inclusion of LRN layers boosted classification accuracy, although such layers have more recently fallen out of favor with the ML community [168]. The first LRN layer in the MsCNN is represented schematically in Figure 4.9d. Note that the LRN layer does not

⁶¹ The inclusion of positive-feedback operations in CNNs is motivated by the behavior of *neurons* in the human brain, the activations of which are often modeled by a hyperbolic tangent function. It is believed that their inclusion improves CNN training performance by introducing instabilities which allow different sets of *neurons* to become active as their weights are adjusted during training.

⁶² The use of LRN is considered analogous to "lateral inhibition" in the human brain; a process by which a highlyexcited *neuron* suppresses the activations of its surrounding *neurons* [172].

alter the size of the data volume, i.e. the output of the first LRN layer in the MsCNN is of size 55 pixels \times 55 pixels \times 96 pixels.

The dimensionality of a CNN (i.e. the size of the data volume) would increase unsustainably through the depth of the CNN (see the following paragraph) without down-sampling (*pooling*) the *responses* from the lower layers. There are several methods by which down-sampling may be achieved⁶³, but all of them operate spatially, i.e. dimensionality is reduced along the width and height of the data volume without affecting the depth of the volume [168]. In the presented MsCNN, down-sampling is accomplished via a Max Pooling layer (Figure 4.9e) [153], [166]. Max Pooling operates by only passing the maximum *response* within a given window on to the next layer. For example, the window size of the first Max Pooling layer of the MsCNN is 3 pixels × 3 pixels [153], [166], therefore only the maximum of the nine *responses* within a window is passed on to the next layer [168]. Interestingly, while *pooling* windows are traditionally non-overlapping, all of the Max Pooling layers in AlexNet utilize windows of size 3 pixels × 3 pixels and a stride of two and therefore operate on overlapping regions [153]. In addition to reducing the dimensionality of the CNN, *pooling* operations have also been shown to mitigate overfitting⁶⁴ [153], [168].

Following the input layer, CONV layer, ReLU layer, LRN layer, and Max Pooling layer, the data volume is once again convolved with a set of *kernels* and the *responses* are stored in a second CONV layer. Notably, while the first CONV layer extracts low-level *features* such as

⁶³ For example, the Average Pooling method takes the mean of the *responses* within the pooling window [168].
⁶⁴ The dimensionality of a CNN is analogous to the degrees of freedom available to a fitted model. Therefore it is not surprising that high-dimensional CNNs may be sensitive to overfitting [165]. Down-sampling via Max Pooling reduces the degrees of freedom of the model and slightly perturbs the training data in the spatial dimensions.

blobs, edges, and lines (Table 4.2), the second CONV layer extracts higher-level *features* [152]. For example, the second CONV layer's analysis of the data volume (containing the *responses* from vertical and horizontal lines) may allow for the detection of intersections of vertical and horizontal lines, e.g. corners. This process is repeated through the depth of the MsCNN for a total of five CONV layers (including the initial CONV layer) with each CONV layer extracting higher and higher level *features*⁶⁵. While each of the four subsequent CONV layers in AlexNet and the presented MsCNN are followed by ReLU layers, only some of the ReLU layers are followed by an LRN layer or a Max Pooling layer [153], [166].

After the final CONV layer and associated ReLU layer, a fully connected (FC) layer (Figure 4.9f) is constructed of size 1 pixel \times 1 pixel \times 4096 pixels. An FC layer is equivalent to a CONV layer in which each *kernel* has a spatial size (width and height) equal to that of the input data volume⁶⁶; therefore each convolution operation produces a single *response* [168]. Finally, a dropout layer (Figure 4.9g) randomly assigns a subset⁶⁷ of the elements within the prior FC layer a value of zero [173]. While not an immediately intuitive operation, the incorporation of dropout layers has been experimentally shown to improve the robustness⁶⁸ of the learned *features* [153]. It is worth noting the CNN architecture originally described by Krizhevsky et al.

⁶⁵ Note that the higher-level *features* are likely to be more specific to the type of data and classifications used during training than the lower-level *features* [152]. In other words, while the *kernels* learned for the first CONV layer are highly similar between trained CNNs, the *kernels* learned for the fifth CONV layer may be quite different between a CNN trained to classify dog breeds and a CNN trained to classify tumors.

⁶⁶ The nomenclature of the FC layer emerges from the concept that each *neuron* (element in the 1 pixel \times 1 pixel \times *n* pixels vector) is connected to everyone *neuron* in the preceding layer. In other words, the *receptive field* of each *neuron* in the FC layer spans the entire input volume and therefore the *response* of each *neuron* is dependent upon the distribution of the *responses* of every *neuron* in the input volume.

⁶⁷ In this case the set of elements is of size 2048 and is composed of unique values $\subseteq \mathbb{Z}[1, 4096]$.

⁶⁸ Increased *feature* robustness is achieved because the dropout operation disrupts the co-dependences between *neurons* thereby encouraging the learning of *neurons* (*kernels*) which do not rely upon the *responses* of other *neurons*.

[153] implemented FC layers slightly differently than explained above in order to facilitate parallelization of the CNN across multiple GPUs. For the sake of clarity, this parallelization has been implicitly ignored throughout this subsection.

4.4.5 Patch-wise Classification

The 1 pixel \times 1 pixel \times 4096 pixels data volume which exists following the final dropout layer (Figure 4.9h) represents the position of the input data (i.e. the current *multi-scale patch*) in high dimensional *feature* space. In other words, this data volume can be thought of as describing the location of the input data along 4,096 axes, with each axis corresponding to a *feature* which has been learned during the training of the AlexNet CNN. Assuming that robust *features* have been learned, it is expected that visually similar input data will have coordinates in common regions of *feature* space.

The final layers of the MsCNN convert this coordinate in *feature* space into the powder bed anomaly classifications. This conversion is achieved through the application of another FC layer (Figure 4.9i) with a depth equal to the number of classification categories; in the case of the MsCNN, the final FC layer is of size 1 pixel \times 1 pixel \times 6 pixels. Recall that each element in a FC layer stores the *response* of a single *kernel* convolved with the entirety of the input volume. Therefore, ideally, the *kernel* learned for the first element of the final FC layer will only produce a high *response* when convolved with a *feature* space coordinate representing an input data belonging to the first classification category. For example, the learned *kernel* corresponding to the third element of the final FC layer nominally produces high *responses* when the input *multi*- *scale patch* contains *recoater streaking* but nominally produces low *responses* when *recoater streaking* is not present in the input *multi-scale patch*.

Finally, a softmax layer (Figure 4.9j) [174] is used to convert the *responses* (of arbitrary magnitude) stored in the final FC layer into a pseudo-probability for each classification category. In the implementation used in this work, the classification category with the largest *response* will have the highest pseudo-probability; the sum of the pseudo-probabilities across the six anomaly classes is always equal to unity. Classification of the input *multi-scale patch* is now trivial. Two classification schemes are explored in this work: (1) Direct classification based solely on the most probable anomaly class (this is shown in Figure 4.9 and is used throughout this chapter and Chapters 3 and 5 unless otherwise specified). (2) Classification based on the two most probable anomaly classes (this is discussed further in Section 4.6.6). Note that unlike the BoW methodology, anomaly classification via the MsCNN methodology operates using a single *patch* (Level 1) of size 25 pixels \times 25 pixels which corresponds to a minimum classifiable powder bed anomaly size of 7.0 mm \times 7.0 mm (refer to Section 4.2.3).

4.4.6 Training

The previous three subsections describe the architecture of the MsCNN and the operations performed on the input data during classification. This subsection is intended to provide a brief overview of the training process for the original AlexNet CNN as well as the application of transfer learning used to convert it to an MsCNN capable of classifying powder bed anomalies. Only the training parameters used by the author for transfer learning are provided below; refer to [153] for a more complete discussion regarding the training of the AlexNet CNN.

116

CNN training operates using a process known as backpropagation [175]. Initially, all of the *weights* of all of the *kernels* throughout the depth of a CNN are randomized [175]. While not previously discussed explicitly, *weights* are simply the element-wise values composing a *filter* or *kernel*. Example *weights* are shown for the vertical edge *filter* in Figure 4.8. During the "forward pass" stage of backpropagation, the training data are passed through the depth of the CNN; because the *kernel weights* are initially randomized, the classification performance will initially be extremely poor.

Fortunately, because the training data are labeled with ground-truth classifications (Section 4.4.2), the performance of the untrained CNN can be quantified. Recall from the previous subsection that the output of the softmax layer is a vector of size $1 \times 1 \times 6$ with values of the set $\mathbb{R}[0,1]$, the sum of which is unity. Therefore, the nominal softmax output for an input *multi-scale patch* containing *recoater streaking* is [0,0,1,0,0,0]. The error (*E*) between this nominal softmax output and the softmax output of the untrained CNN can be defined using a variety of energy functions⁶⁹. The error (or loss) is a function of the set of *weights* (Ω); Figure 4.10 shows a graphical representation of a loss function in the simplified case of only two *weights* (ω_1, ω_2).

⁶⁹ Perhaps the simplest energy function is Mean Square Error (MSE), though the MATLAB implementation of AlexNet calculates the error using a "cross entropy function for mutually exclusive classes" [177].



Figure 4.10: A generalized loss function for the case of $\Omega[\omega_1, \omega_2]$ represented by a concave surface in 3D space. Refer to the proceeding and following paragraphs for discussion and variable definitions. This figure is a modified version of the figure presented by Deshpande [175].

The current, non-optimized *weights* are shown in Figure 4.10 as point Ω_i on the surface of the loss function. As the goal is to reduce the classification error, it is desirable to adjust the *weights* in the direction opposite to the gradient of the loss function (∇E) – the calculation of the gradient is considered the "backward pass" stage of the backpropagation process [175]. Both AlexNet and the MsCNN utilize a method, defined in (4.3), known as Stochastic Gradient Descent with a Momentum term (SGDM) to calculate the *weight* adjustment [153], [166]. Each update to the *weights* is known as an iteration (i) and is often referred to as the "parameter update" stage of the backpropagation process [175].

$$\Omega_{i+1} = \Omega_i - \eta \nabla E(\Omega_i) + \gamma (\Omega_i - \Omega_{i-1})$$
(4.3)

Where Ω_{i+1} is the updated set of *weights*, Ω_i is the current set of *weights*, η is the learning rate, $\nabla E(\Omega_i)$ is the gradient of the loss function evaluated at the current set of *weights*, γ is the momentum coefficient, and Ω_{i-1} is the set of *weights* during the previous iteration.

While the direction to adjust the *weights* is derived from the gradient of the loss function, the magnitude of the *weight* adjustment in the derived direction is defined a priori and is known as the "learning rate" (η). A higher learning rate has the potential to increase the learning speed⁷⁰ but if the learning rate is too high, learning may become unstable and convergence to a local minimum may not occur [175]. It is also common for the gradient to oscillate between iterations; one approach to mitigate this oscillation is to incorporate a "momentum" term which biases the calculated gradient in the direction of the gradient calculated during the previous iteration. The impact the momentum term on the calculation of the current weight adjustment is controlled by the coefficient γ . In traditional GD, the loss function is defined for the entirety of the training dataset. While this approach can produce high classification accuracies, it is too computationally expensive to be used for backpropagation through the depth of a CNN [176]. For this reason, AlexNet utilizes SGD which defines the loss function only over a subset of the training dataset [153], [166]. Each subset of the training dataset is known as a "mini-batch" and is randomly (hence the "stochastic" nomenclature) delineated at runtime [177]. Each time convergence is achieved for the set of mini-batches covering the entire dataset, the entire backpropagation process is repeated; each repetition is referred to as an "epoch." During subsequent epochs, the initial weights are those learned during this previous epoch, i.e. they are not randomized.

⁷⁰ One method of increasing the training speed without risking divergence is to start the training process with an initial learning rate that is relatively low. Then, because the initial *weights* improve during the training process (i.e. the initial guesses lie increasingly close to the local minimum), the learning rate can be increased (from its initial value) at the start of subsequent epochs without causing divergence; this approach is often referred to as "scheduling" the learning rate [176], [177].

During the training of a full CNN such as presented by Krizhevsky et al. [153] all of the *weights* are initialized randomly and backpropagation is applied through the depth of the CNN. As previously mentioned, such a task is computationally expensive; fortunately, transfer learning allowed the presented MsCNN to begin its training with many of the lower-level *weights* already learned. Specifically, only the final FC layer of size 1 pixel \times 1 pixel \times 6 pixels (Figure 4.9i) and the softmax layer (Figure 4.9j) are trained on the training database of *multiscale patches*; all of the *weights* throughout the rest of the depth of the MsCNN remain identical to the *weights* of the pre-trained AlexNet CNN. For training of the final two layers a constant (i.e. unscheduled) learning rate⁷¹ of 0.001 was used and a total of 20 epochs were executed. A momentum coefficient of 0.9 was used and each mini-batch⁷² contained 64 *multiscale patches* from the training database. All other training parameters were set to the defaults listed in [177].

Finally, it should be noted that during the described training process, only the *kernel weights* are learned. In other words, the architecture of the CNN remains static and is not automatically optimized. During the CNN design process a human programmer manually modifies the CNN architecture in order to achieve improved *validation* performance (Section 4.5.2). Common architecture adjustments are often referred to as hyperparameters and include the overall depth of the CNN, the types of layers, the size and stride of the *kernels*, and even the pre-processing applied to the input data [152], [153]. Because transfer learning was only

⁷¹ An appropriate learning rate was determined by starting with a rate of 0.1 and reducing the rate by factors of ten until stable convergence during training was achieved.

⁷² The mini-batch size was limited by the amount of RAM onboard the GPU as discussed in Section 4.5.4.

applied to the final two layers of the AlexNet CNN, hyperparameter tuning was not explored in this work. An excellent discussion of this topic is presented in [153].

4.4.7 Layer-wise Classification

In contrast to the BoW methodology, only two contextual heuristics are implemented. The observant reader may have noticed that while there are seven classification categories presented in Section 4.2.4, the MsCNN only produces pseudo-probabilities for six classification categories. Due to the visual similarity between some *patches* containing *debris* and some *patches* containing *part damage*, the two categories were combined into the *disturbance* classification for the purposes of MsCNN training and classification. Final differentiation between the *debris* and *part damage* anomaly types is implemented layer-wise and is based solely on the location of the part geometry cross-sections in the current layer. Specifically, if a pixel not lying on top of a part (as defined by the extracted part geometry information) is classified by the MsCNN as *disturbance*, it will be re-labeled as *debris*. Conversely, if a pixel lying on top of a part is classified by the MsCNN as *disturbance*, it will be re-labeled as *super-elevation*, it will be re-labeled as *debris* (see Section 4.2.4 for justification).

Finally, the layer-wise anomaly classifications for an entire build are converted from a 3D point cloud to a compressed data format in which a number quadruplet encodes the type of anomaly and its location in 3D space while the locations of the *okay* pixels are not explicitly stored. Because the anomaly classification information is always used in a layer-wise fashion,

the speed of data decompression is improved by independently storing the location of the first anomalous pixel in each layer within the compressed data structure.

4.5 Performance of the ML Methodologies

4.5.1 Confusion Matrices

At this point, the reader is encouraged to recall the high-level discussion of ML provided in Section 1.4. It is common to evaluate ML algorithms using a metric known as a confusion matrix [178]. Fundamentally, a confusion matrix compares a ML algorithm's classifications to the ground truth classifications made by a human. In all implementations, the data used to generate a *confusion matrix* must be separate from the data used to train the ML algorithm. Traditionally, the entire available dataset (i.e. the training database in Sections 4.3.2 and 4.4.2) is divided into three subsets known as training, validation, and testing datasets [179]. During the training process, the ML model is fit to the *training* data set. The performance of the model can then be evaluated using the validation dataset and the human programmer may decide to modify the design of the model based on these results. Common modifications may include alterations to the input data format (e.g. *patch* padding or lighting calibrations), increasing the amount of training data, or tuning hyperparameters such as those described in Section 4.4.6. Once the design of the ML algorithm and any accompanying methodology is complete, the true performance can be estimated using the *testing* dataset. The *testing* dataset also serves as a final check that the model has not been over-fit to the *training* data [165].

Due to the layer-wise (as opposed to patch-wise) nature of the final anomaly classifications produced by the BoW methodology, creation of a meaningful *validation* data set is not possible.

122

In other words, because the contextual heuristics (Section 4.3.6) play a key role in anomaly classification, the performance of the BoW ML model cannot be independently validated. Similarly, the *testing* performance must be evaluated on a layer-wise basis. Layer-wise ground truths were generated via the manual classification of hundreds of thousands of pixels across twenty representative powder bed images (provided in Appendix D). These ground truth classifications are then compared, pixel-wise, across the entire layer with the classifications produced by the BoW methodology. In order to enable direct comparison between the *testing* performances of the BoW and MsCNN methodologies, the same layer-wise ground truths and evaluation procedure are used for both methodologies. It should be noted that because ground truth anomaly classification across an entire powder bed often necessitates judgment calls (on the part of the human) for ambiguous regions (which would typically not be included in a *training* or *validation* dataset), the reported *testing* performance is expected to be lower than the *validation* performance as well as the *testing* performances reported in the literature for more "typical" CV applications of ML algorithms.

4.5.2 Validation Accuracy

The values along the left-to-right diagonal of the *confusion matrices* (Figures 4.11 and 4.12) represent the percentage of *patches* for which the ML algorithms make classifications identical to those of the human. In other words, if an algorithm evaluates the *validation* dataset exactly the same as the human, all of the diagonal values will be 100%. All of the diagonal values less than 100% represent the degree to which the ML algorithm has been "confused" with regards to classification of specific anomaly. The *validation* dataset in this work is produced via a

random partitioning of the training database (Section 4.4.2) such that 90% of the *patches* within the training database are used for *training* and the remaining 10% of *patches* are used for *validation*.

The *validation* accuracies are reported for both the MsCNN described in Section 4.4 as well as a CNN which does not utilize multi-scale input data. Specifically, the Level 1 *patch* is duplicated across all three input channels (Figure 4.9). The CNN is otherwise structurally identical to the MsCNN. Curiously, the *validation* performances are not significantly different between the CNN and MsCNN algorithms; the *testing* accuracies, and a more expansive discussion, are reported in the following subsection. As noted in Section 4.4.1, the use of transfer learning precluded the possibility of tuning the MsCNN hyperparameters based on the *validation* performance results.



Figure 4.11: A confusion matrix developed from the validation dataset and reporting the percentage of *patches* classified correctly by the **CNN algorithm** for each anomaly type. The absolute numbers of *patches* classified by the human as each anomaly type are shown in parentheses on the vertical axis.

Figure 4.12: A confusion matrix developed from the validation dataset and reporting the percentage of *patches* classified correctly by the **MsCNN algorithm** for each anomaly type.

To determine the dependence of the MsCNN classification accuracy on the size of the training database, Figure 4.13 reports the *validation* performance for the MsCNN trained using the entirety of the available training database as well as for the MsCNN trained using only 10% of the available training database. To ensure robust measures of *validation* performance, k-fold cross-validation [179], [180, p. 78] was performed using ten segments⁷³ to generate "average" *validation* classification accuracies for each anomaly type. Overall, the *validation* performance

⁷³ Specifically, k-fold cross-validation was implemented by randomly partitioning the training database into 10 equally-sized segments. During the first k-fold iteration, *training* was performed using the *patches* within segments #2 - #10 while *validation* was performed using the *patches* within segment #1. During the second k-fold iteration, *training* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #1 and #3 - #10 while *validation* was performed using the *patches* within segments #2. A total of 10 iterations were performed such that the segments were used for *training* and *validation* in all possible combinations.

remained remarkably robust despite an order of magnitude decrease in the size of the training database. The most significant performance drops were observed for *recoater streaking* (the anomaly with the fewest available *patches* in the training database), *disturbances*, and *super-elevation* (arguably the anomalies with the highest degree of variability in their appearance). The performance drops (in absolute percentage points) are 7%, 13%, and 3%, respectively.



Figure 4.13: Sensitivity of the MsCNN algorithm to training database size. The reported accuracy percentages are the average generated by the described k-fold cross-validation procedure. The error bars represent 95% confidence intervals about the mean.

4.5.3 Testing Accuracy

As in the previous section, the values along the left-to-right diagonal of the *confusion matrices* (Figures 4.14 – 4.19) represent the percentage of pixels for which the ML methodologies make classifications identical to those of the human. In other words, if a methodology evaluates the *testing* dataset exactly the same as the human, all of the diagonal

values will be 100%. All of the diagonal values less than 100% represent the degree to which the ML methodology has been "confused" with regards to classification of specific anomaly.

The percentages of "anomalies classified correctly" are reported in Figures 4.14 – 4.16 for the BoW, CNN, and MsCNN methodologies, respectively. For example, 94.2% of the pixels classified by a human as part damage are correctly classified by the MsCNN methodology as part damage. Essentially, these three confusion matrices represent how well the three methodologies classify each anomaly. It is immediately apparent that the CNN methodology performs drastically worse than either the BoW or MsCNN methodologies. This is to be expected as, unlike the other two methodologies, no data about the broader region surrounding a *patch* are incorporated in the classification process. It is worth noting that the low testing performance of the CNN methodology lies in stark contrast its relatively high validation performance. As recognized in the literature, such a discrepancy is often an indication that the ML model has been over-fit to the *training* data [179]. This discrepancy also suggests that there is a degree of patch variation present in the testing data which is not wellrepresented within the *training* or *validation* data sets. The performance of the CNN was only lightly investigated by the author as the primary focus of this chapter is on the BoW and MsCNN methodologies.

Due to their small size (relative to the camera resolution) *recoater hopping* and *recoater streaking* were among the most difficult anomalies for both the BoW (62.7%, 39.5%) and the MsCNN (72.7%, 76.9%) methodologies to classify. For the MsCNN methodology, the classification accuracy for *super-elevation* (73.0%) was also low relative to the other classification accuracies. However, in contrast to the BoW methodology, the majority of the

127

MsCNN miss-classifications of *super-elevation* are classified as *part damage* (16.4% out of 27.0%) which often occurs concurrently with *super-elevation*. Similarly, the majority of MsCNN miss-classifications of *recoater streaking* are classified as *debris* (13.5% out of 23.1%) which is often co-located with *recoater streaking*. The lowest classification accuracies are 39.5%, 0.0%, and 72.7% for the BoW, CNN, and MsCNN methodologies, respectively. The performance of the MsCNN is superior for all of the anomaly types.

Human (Ground Truth) Anomaly Classification



developed from the testing dataset and reporting the percentage of pixels classified BoW correctly by the methodology for each anomaly type.

developed from testing the dataset and reporting the percentage of pixels classified correctly by the CNN methodology for each anomaly type.

Figure 4.16: A confusion matrix developed from the testing dataset and reporting the percentage of pixels classified correctly by the MsCNN methodology for each anomaly type.

Incomplete Spreading

81.9%

Part Damage

For clarity purposes, the labels on the vertical axis have been neglected. The anomaly types and the corresponding number of pixels classified by the human are as follows, from top to bottom: okay (10,651,250), recoater hopping (478,750), recoater streaking (40,625), debris (262,244), super-elevation (130,874), part damage (47,131), and incomplete spreading (263,125).

⁷⁴ The *testing* accuracies reported in this figure are lower than those reported by the author in Scime et al. [54, Fig. 5]. Note that the while the powder bed images used for *testing* are identical in both works, in this thesis the *testing* images have been more fully classified by the human in order to decrease potential bias. As a result, more ambiguous regions of the powder bed have been included the in the performance characterization, resulting in a more aggressive *testing* performance metric than used previously.

Alternatively, the percentages of "anomaly classifications that are correct" are reported in Figures 4.17 – 4.19 for the BoW, CNN, and MsCNN methodologies, respectively. For example, 63.7% of the pixels classified by the MsCNN methodology as *part damage* are truly *part damage* (as defined by the human classification). In other words, while the MsCNN is likely to catch 94.2% of *part damage* regions, 36.8% of its *part damage* classifications are incorrect (30.3% out of those 36.8% are instead instances of *super-elevation* as defined by the human classification). Essentially, these *confusion matrices* represent how well the ML methodologies avoid false-classifications⁷⁵ of each anomaly. As in the previous *confusion matrices*, the CNN methodology performed dramatically worse than the other methodologies while the MsCNN methodology performed the best in the case of all anomalies except for *part damage*, for which the BoW mythology avoided false classifications at a notably higher rate (74.6% to 63.7%).

⁷⁵ Certain modifications to the ML methodology or the composition of the training database may improve the classification accuracy of a targeted anomaly type to the detriment of the classification of other anomaly types. In such a situation, the diagonal percentage in Figure 4.14 (or 4.15 or 4.16) corresponding to the targeted anomaly would increase, while the corresponding diagonal percentage in Figure 4.17 (or 4.18 or 4.19) would decrease, i.e. the rate of false-classifications for that anomaly would increase.



Figure 4.17: A confusion matrix^{7°} developed from the *testing* dataset and reporting the percentage of pixel classifications made by the **BoW methodology** that are correct.

Figure 4.18: A *confusion matrix*⁷⁷ developed from the *testing* dataset and reporting the percentage of pixel classifications made by the **CNN methodology** that are correct.

Figure 4.19: A *confusion matrix* developed from the *testing* dataset and reporting the percentage of pixel classifications made by the **MsCNN methodology** that are correct.

For clarity purposes, the labels on the vertical axis have been neglected. The anomaly types and the corresponding number of pixels classified by the human are as follows, from top to bottom: *okay* (10,651,250), *recoater hopping* (478,750), *recoater streaking* (40,625), *debris* (262,244), *super-elevation* (130,874), *part damage* (47,131), and *incomplete spreading* (263,125).

⁷⁶ The *testing* accuracies reported in this figure are lower than those reported by the author in [54, Fig. 6]. Note that the while the powder bed images used for *testing* are identical in both works, in this thesis the *testing* images have been more fully classified by the human in order to decrease potential bias. As a result, more ambiguous regions of the powder bed have been included the in the performance characterization, resulting in a more aggressive *testing* performance metric than used previously.

⁷⁷ The column of undefined (NaN) values is the result the CNN's failure to classify any instances of *recoater streaking* in the *testing* dataset.

A common metric for ML methodology *testing* performance is the overall guessing accuracy, which in this case would be defined as the percentage of pixels that the ML methodologies classified the same as the human. However, because the overwhelming majority (90%) of the pixels in the *testing* dataset are *okay*, the author does not consider this metric to be fully informative in this situation. Therefore, two additional performance metrics are also considered: The anomaly detection accuracy, i.e. with what accuracy is the presence of an anomaly (i.e. a non-*okay* pixel) detected by the ML methodology. And the total classification accuracy among the anomalies, i.e. if a pixel is correctly classified as anomalous (i.e. non-*okay*), with what accuracy is the pixel correctly classified by the ML methodology. These three metrics are summarized in Table 4.3 for the two primary ML methodologies.

Table 4.3: Overall testing	data cl	assification	accuracy	/ metrics
----------------------------	---------	--------------	----------	-----------

Accuracy Metric	BoW	MsCNN
Overall Accuracy	95%	97%
Anomaly Detection Accuracy	73%	85%
Accuracy among the Anomalies	88%	93%

4.5.4 Computational Burden

All of the following computational calculations were performed on a system with the specifications listed in Table 4.4. Processing times on a per layer basis for both of the primary ML methodologies as well as the ancillary operations are reported in Table 4.5. The typical training times for the final BoW and MsCNN algorithms are 15 hours and 1 hour, respectively. Due to the design of the compression algorithm, the number of analyzed layers does not correspond directly to the memory burden of the build analysis. The largest analyzed build by number of layers is 5,452 layers tall and requires 37 MB on the hard disk and 2 GB when loaded

into RAM, while the largest analyzed build by memory requirement is only 1,067 layers tall but requires 96 MB on the hard disk and 2.7 GB when loaded into RAM.

Table 4.4: Computer system specifications.

Specification	Value	
Operating System	Microsoft [®] Windows [®] 10	
System RAM	HyperX [®] 4 $ imes$ 8 GB (DDR4)	
Processor Type (CPU)	Intel [®] i7-6700K (4 cores)	
Processor Speed (CPU)	4.00 GHZ	
Graphics Card (GPU)	NVIDIA [®] GeForce GTX 950	
Dedicated/On-Board RAM (GPU)	2 GB	

Table 4.5: Approximate computation time for selected layer-wise operations.

Layer-wise Operation	Computation Time (seconds/layer)
EOSPRINT Screenshot Capture	5
Part Geometry Extraction	0.2
3D Data Decompression	0.02
BoW Classification	4
MsCNN Classification	7

The computation speed of the BoW algorithm could be further increased via explicit parallelization of the layer analysis process – while some of the utilized MATLAB functions may be implicitly parallelized, the implementation described in this work only explicitly uses one CPU core to perform the calculations. As the patch-wise classifications are independent from each other, threading these calculations⁷⁸ would be trivial. The computation speed of the MsCNN algorithm could be further increased though the use of a research-grade GPU or even an array of multiple graphics cards. Of note, the current system's limiting specification appears to be the RAM on-board the Graphics Card as opposed available GPU clock cycles. More on-

⁷⁸ This could be implemented by each available core performing patch-wise classification on a different set of *patches* from the same layer.

board RAM would allow for larger mini-batch sizes which would allow for a higher utilization of the GPU itself⁷⁹.

Powder spreading in the EOS M290 requires approximately six seconds and exposure of a typical layer requires on the order of several minutes. Therefore, the current implementations of both the BoW and the MsCNN algorithms operate fast enough to be considered "real-time." The use of powder bed anomaly classification in a real-time environment is discussed in Section 5.3 and the benefits of increased computation speeds are discussed in Section 8.3.

4.5.5 Comparison of the Machine Learning Methodologies

The MsCNN methodology was determined by the author to be superior to the BoW methodology for the reasons enumerated below. Therefore, the MsCNN methodology is the sole ML methodology used for the remainder of this chapter as well as the entirety of Chapter 5; for convenience it is henceforth often referred to as "the MsCNN."

- The confusion matrices reported earlier in this section clearly demonstrate that the MsCNN methodology provides far superior anomaly classification performance, relative the BoW methodology, in nearly all cases.
- 2. The MsCNN methodology's uniform anomaly classification resolution of 25 pixels \times 25 pixels (defined by the *patch* size, Section 4.4.2) is comparable to the finest resolution offered by the BoW methodology (20 pixels \times 20 pixels) and it is substantially finer than the BoW methodology's 100 pixels \times 100 pixels *incomplete spreading* resolution.

⁷⁹ The mini-batch size during classification is limited to 256 *multi-scale patches* of the powder bed image due to the RAM available onboard the GPU. Because there are a total of 1,296 *multi-scale patches* in each powder bed image, a mini-batch size of 1,296 would be the most computationally efficient for the presented implementation.

- The pseudo-probabilistic classifications produced by the MsCNN's softmax layer (Section 4.4.5 and 4.6.6) have the potential to provide an operator or a feedback control algorithm with valuable information not provided by the BoW methodology.
- 4. The lack of human-designed components such as the *filter bank* and the contextual heuristics endows the MsCNN methodology with greater flexibility during retraining (e.g. if the camera resolution is altered or a different L-PBF machine is used) than the BoW methodology.

Note that while layer-wise anomaly classification via the MsCNN methodology is slightly slower than via the BoW methodology, its computation speed can be dramatically improved as described in the previous subsection. Finally, the respective training times of the two ML methodologies were not considered in the above decision process as they are highly dependent upon the exact training parameters used (Sections 4.3.4 and 4.4.6).

4.5.6 Portability between Material Systems and L-PBF Machines

The MsCNN methodology has been used to analyze builds using multiple material systems including: AlSi10Mg, bronze, Inconel 625, Inconel 718 (two Powder Systems), stainless steel 316L, stainless steel 17-4 PH, and Ti-6Al-4V (four Powder Systems). Performance of the MsCNN is robust across all of the material systems and there is no evident need for material system-specific training databases. The training database includes data from builds using most, but not all, of the above material systems. Additional details about the use of bronze (the most visually-distinct material system) and AlSi10Mg (the material system with the highest background level

of anomaly detections) are provided in Section 5.2.4. The use of the MsCNN methodology to correlate powder bed anomalies with powder particle size is reported in Section 3.3.2.

While the vast majority of the 51 builds analyzed by the MsCNN were performed on the EOS M290 machine at CMU's NextManufacturing Center, two builds performed on the EOS M290 at Eaton's® Additive Manufacturing Center of Excellence facility in Southfield, Michigan were also analyzed. No structural changes to the MsCNN or retraining were performed; successful analysis of the builds only required a machine-specific calibration, to compensate for differing camera positons and lighting conditions, as outlined in Section 4.2.3. Additionally, powder bed images captured during a build performed on an SLM 280 (SLM Solutions GmbH, Germany) L-PBF machine at Arconic's Production Center in Austin, Texas were also successfully analyzed. Note that in this case several manual lighting adjustments were required in addition to the standard calibration procedure. The author suspects that such manual adjustments would not be necessary if the MsCNN were retrained with data specific to SLM 280 machines. Overall, over 70,000 powder bed images have been analyzed using the MsCNN.

Finally, it is worth noting that the entire MsCNN methodology is sufficiently robust to be operated as a set of standalone executables. Indeed, the software pipeline has been designed such that it can be installed on a computer without MATLAB, although the free MATLAB runtime environment is still required [181]. Internal tests have demonstrated that the software can allow a user with minimal knowledge of L-PBF, CV, or ML to perform the necessary calibrations, retrain the MsCNN, analyze a build, and interpret the analysis results. Figure 4.20 shows a screenshot of the user interface for a portion of the standalone software package.

136



Figure 4.20: A screenshot of the Graphical User Interface for a portion of the installed powder bed anomaly classification software package.

4.6 Results

4.6.1 Case Study Overview

In this chapter, a single case study is used to illustrate the capabilities of the MsCNN methodology. Additional case studies are explored in Chapter 5. The build examined in this section was performed in December of 2016 for Dr. Samikshya Subedi and Prof. Anthony Rollett of CMU's Materials Science and Engineering Department under a US Department of Energy grant (DE-FE0024064) and in collaboration with Oregon State University and Prof. Vinod Narayanan at the University of California Davis. Figure 4.21 shows a CAD rendering of the heat exchanger built for this project out of Inconel 718 (In718) using 20 µm thick layers.



Figure 4.21: A CAD rendering of one of the heat exchangers included in the build discussed in this section.

4.6.2 Analysis of a Single Layer

Figure 4.23 shows the raw powder bed image taken at layer 2,709 of the heat exchanger build. The approximate location of this layer is indicated by the solid blue line in Figure 4.21. The heat exchanger of interest is shown within the solid blue bounding box. Note that an identical heat exchanger, within the dotted blue bounding box, was deactivated after layer 1,894 – an event discussed in the following subsection. A number of different powder bed anomalies are present, a selection of which have been annotated for clarity. For example, in several locations the heat exchanger and surrounding parts have warped up above the powder bed due to a buildup of residual thermal stresses (*super-elevation*). This distortion has allowed the recoater blade to impact the heat exchanger resulting in vertical (parallel to the *y*-axis) markings characteristic of *recoater hopping*. This impact has also resulted in substantial *part damage* and significant quantities of *debris* in the surrounding powder bed. Evidence of *recoater streaking* due to a combination of damage to the recoater blade and the dragging of debris is also visible. Finally, to conserve In718 powder, the dosing factor (the amount of

powder fetched from the powder dispenser and spread for each layer) was reduced, resulting in *incomplete spreading* toward the left-hand extreme of the powder bed.



Figure 4.22: Layer 2,709 (54.18 mm above the build plate) with selected instances of powder bed anomalies manually annotated. The heat exchanger of interest is inside the solid blue bounding box while the heat exchanger within the dotted blue bounding box was disabled earlier in the build by the operator.

Each layer of a build is analyzed by the MsCNN as described in Section 4.4. Figure 4.23 shows the MsCNN analysis of layer 2,709. The regions of *super-elevation* are highlighted in red and the majority of the *recoater hopping* is highlighted in teal. Instances of *part damage* are highlighted in magenta and many instances of *debris* are highlighted in white. Many of the instances of *recoater streaking* are highlighted in blue, although as suggested by the *confusion matrices* in Section 4.5.3, some of the instances of *recoater streaking* are not detected by the MsCNN. Finally, the regions experiencing *incomplete spreading* are highlighted in yellow; note that as suggested by the *confusion matrices* in Section 4.5.3, *incomplete spreading* is occasionally miss-classified as *debris*.


Figure 4.23: Layer 2,709 (54.18 mm above the build plate) with the powder bed anomalies classified by the MsCNN. The green pixels indicate the CAD outline of the parts at this layer.

4.6.3 Global and Local Build Reports

After each layer has been analyzed, the percentage of each anomaly classification in each layer can be displayed as a function of build height in a *build report*. Figure 4.24 shows a simplified (only containing two anomaly types) *global build report*. Note the anomaly spikes around layer 2,709 (Figure 4.23) and the spike at layer 1,894 that corresponds to the operator intervening in the build process and deactivating the second heat exchanger (dotted blue bounding box in Figure 4.23).



Figure 4.24: A *global build report* showing the number of pixels (as a percentage of part area, based on the CAD geometry) classified as *super-elevation* and *part damage* anomalies at each layer of the build. Note layers 1,894 and 2,709 which are discussed in this and previous subsections.

Local build reports can also be created, showing only the anomaly classifications results within a specified region of the powder bed; this is particularly useful if there are multiple parts in a build. Figure 4.25 shows a *local build report* covering the region bounded by the solid blue box in Figure 4.22. The lower subplot (green) of Figure 4.25 shows the percentage of pixels within the specified region that lie on top of the heat exchanger itself (based on CAD geometry) as a function of build height. In other words, the lower subplot shows the percentage of the bounded area that was fused in the previous layer. Equivalently, the lower subplot shows the amount of fused material covered by only one layer of powder. Recognize that an increase in this percentage indicates an overhanging region, and a sharp increase implies a largely unsupported overhang. Overhanging regions are discussed in more detail throughout Chapter 5. Note that many of the anomaly detections (layers 2,500 – 3,000) coincide with the relatively rapid geometry changes, including the construction of an overhang, that occur between 50 mm and 60 mm of build height.



Figure 4.25: *Local build report* and vertical part profile for the heat exchanger of interest. The anomaly classifications in the upper subplot are presented as percentages of the area of the specified region.

4.6.4 Cumulative Anomaly Classifications

The cumulative occurrences of anomalies throughout the height of the build can be displayed as *heat maps*; this representation can be effective at highlighting regions of the powder bed (i.e. specific parts) that consistently experienced problems throughout the build height. Figure 4.26 shows a *heat map* of *part damage* anomaly classifications over the course of the entire build. The region that stands out (*part damage* occurred for 4% of layers in this location) is the same region discussed in the following subsection and highlighted in magenta in Figure 4.23. Figure 4.27 shows a *heat map* of *incomplete spreading* classifications over the course of the entire build. As expected, these detections occur near the left-hand extreme of the powder bed.



Figure 4.26: A *heat map* showing the percentage of layers (throughout the build height) in which *part damage* was detected at each pixel. The footprints (CAD geometry boundaries at the first layer) of each part are shown as white outlines.



Figure 4.27: A *heat map* showing the percentage of layers (throughout the build height) in which *incomplete spreading* was detected at each pixel. The footprints (CAD geometry boundaries at the first layer) of each part are shown as white outlines.

4.6.5 3D Renderings

The anomalies classified in each layer can be mapped onto a 3D model of the build or a specific part as shown in Figure 4.28, where the detected regions of *part damage* are indicated in magenta. Figure 4.29 shows the as-built heat exchanger; significant damage to the part is clearly visible (within the red bounding box) in the same locations indicated in the 3D rendering.



Figure 4.28: A 3D rendering of the heat exchanger of interest with the detected regions of *part damage* indicated in magenta.

Figure 4.29: The as-built heat exchanger with the visible defects highlighted by a red bounding box.

4.6.6 Pseudo-Probabilistic Anomaly Classifications

As discussed in Section 4.4.7, the final layer of the MsCNN utilizes a softmax function. That is, a pseudo-probability (between 0% and 100%) is assigned to each anomaly type. The classifications reported in Figure 4.23 are based only on the most probable anomaly type – essentially, the additional information provided by the softmax function is ignored. Conversely, the classifications shown in Figure 4.30 incorporate the probability information. Specifically, if the most probable anomaly type for a given *patch* has a pseudo-probability less than 75%, then both the most probable and the second most probable classifications are displayed simultaneously. In this situation, the *patch* is highlighted by alternating vertical stripes, the colors of which indicate the two most probable anomaly classifications.

The visualization of the classification pseudo-probabilities demonstrates the types of situations which the MsCNN finds most ambiguous. For example, in some cases the MsCNN

finds the distinction between *debris* and either *incomplete spreading* or *recoater streaking* to be ambiguous. Additionally, some of the less severe cases of *recoater hopping* appear similar to *okay* regions. Interestingly, the vast majority of the anomaly classifications made for this powder bed image have a pseudo-probability of greater than 75%, that is, the MsCNN is relatively certain of its classification of most *patches*.



Figure 4.30: Layer 2,709 (54.18 mm above the build plate) with the powder bed anomalies classified by the MsCNN. Anomaly classifications with a pseudo-probability less than 75% are shown as a *patch* with hatched coloring. Specifically, within those *patches* alternating vertical stripes are used to indicate the highest and second highest probably anomaly classification. The green pixels indicate the CAD geometry outline of the parts at this layer.

4.6.7 Additional Analysis Visualization Modalities

The anomaly classifications produced by the MsCNN can be viewed in several additional formats including layer-by-layer time-lapses of the build as well as movies showing the layer-by-layer 3D rendering of the build – thereby allowing the observation of anomalies internal to the part geometry. Also of interest, a Fourier frequency analysis can be applied to a *global* or *local build report*. For example, such an analysis of the *recoater hopping* classifications in Figure 4.25 results in an anomaly recurrence period of 250 layers; this is correlated with the periodic geometry changes (part profile subplot) also shown in Figure 4.25. These geometry changes are related to the design of the internal cooling channels within the heat exchanger. Other, more illuminating, examples of frequency analyses are presented in Sections 3.3.2 and 5.2.4 where it enables the correlation of powder bed anomalies with powder particle size, laser scan strategy, and layer-wise energy density.

4.7 Discussion and Summary

In this chapter, autonomous powder bed anomaly detection and classification is achieved through the use of contemporary Machine Learning and Computer Vision techniques. Layerwise powder bed images are captured throughout each build using a relatively low-resolution (1 MP) camera. Notably, no modifications were made to the camera or lighting system provided by the L-PBF machine manufacturer (EOS GmbH). Successful use of the images produced by the manufacturer-provided system required the development a robust, machine-specific calibration process. Six different millimeter-scale powder bed anomaly types are described – while several of these anomaly types have been reported in prior literature, to the author's knowledge, this thesis provides the most comprehensive accounting of L-PBF powder bed anomalies. In order to improve the utility of the autonomous anomaly classifications, part geometry and build layout information are extracted layer-wise from the EOSPRINT software environment using a custom CV algorithm.

Two anomaly classification methodologies are presented, each relying on a different ML algorithm. The first ML algorithm is an application of the well-established Bag of Words approach. While relatively easy to implement, the BoW methodology relies on human-created *feature* extraction tools and requires human-defined contextual heuristics to achieve reasonable classification accuracies for several of the anomaly types. The second ML algorithm uses transfer learning to apply the AlexNet CNN to the powder bed anomaly classification problem. Modern CNNs were developed more recently than the BoW approach and learn their own "optimal" *feature* extraction tools. Furthermore, through the use of *multi-scale patches* the need for contextual heuristics was eliminated.

Interestingly, the author has found extremely few prior examples of analyzing data structures containing multi-scale information with CNNs. This suggests that the use of MsCNNs has significant unexplored potential both in AM and other fields; Chapter 8 discusses some of the other potential AM applications for MsCNNs. Somewhat surprisingly, the re-trained AlexNet CNN provides high classification performance despite the fact that the *kernels* in the first CONV layer were initially trained to extract useful information from three-channel color images instead of three-channel multi-scale images. The author proposes one possible explanation for this pleasant surprise. As can be seen in [153, Fig. 3], many of the *kernels* in AlexNet's first CONV layer are either grayscale (i.e. channel-agnostic) or only produce high *responses* on

147

channel-defined colors (i.e. red, green, or blue). The grayscale and color-specific *kernels* may reside in relatively distinct regions of the final 4,096 dimensional *feature* space and therefore the six *kernels* in the final FC layer (which are retrained during transfer learning) can "cleanly" operate on only the channel-agnostic region of *feature* space. Importantly, Krizhevsky et al. [153, p. 7] observed that while the *kernels* learned on one GPU routinely exhibited color-agnosticism the *kernels* on the second GPU were routinely color-specific; due to their implementation of GPU parallelization it is conceivable that this differentiation propagated into the final *feature* space.

Validation and testing performances were reported for the BoW, MsCNN, and CNN (the AlexNet CNN trained on only Level 1 patches) methodologies. The final MsCNN has higher testing anomaly classification accuracies than the other two methodologies, particularly for the recoater streaking, debris, super-elevation, and part damage anomaly types. For the final MsCNN, the overall, anomaly detection, and anomaly differentiation accuracies are 97%, 85%, and 93%, respectively. While not explored in this work, overall testing accuracy may be increased if the MsCNN is trained on a balanced dataset, that is, a dataset containing an equal number image patches labeled as each anomaly type. Note that the extremely high mini-batch accuracies observed during the later training epochs suggest that convergence is not currently an issue and that the final learned kernels are not disproportionately biased toward the anomaly classes with the greatest number of examples in the training dataset. While the BoW methodology requires less computation time per layer than the MsCNN methodology (4 seconds vs. 7 seconds) this advantage is system-specific and substantially more powerful GPUs than used in this work are commercially available. The higher classification accuracy, higher

incomplete spreading classification resolution, and increased flexibility of the MsCNN over the BoW methodology demonstrated its superior performance for this application.

The final MsCNN methodology was shown to perform robustly across a wide range of material systems and was able to successfully analyze data captured by an alternate EOS M290 machine as well as an SLM 280 machine. The presented methodology was used to analyze over 70,000 powder bed images, producing almost 100 million patch-wise anomaly classifications and can be operated by external users as a stand-alone software package. A case study was used to demonstrate the utility of the autonomous anomaly classifications for build failure analysis. Additional case studies spanning a wide range of build geometries, material systems, and processing conditions are explored in Chapter 5.

The classification results can be viewed in a variety of formats including anomaly classifications as a function of build height, cumulative heat maps, 3D renderings of the part geometries, and time-lapse videos. As demonstrated in Chapter 3, a frequency analysis of the anomaly detections can provide valuable insight into the interaction between the laser scan strategy, the powder feedstock, and disturbances to the powder bed. While only briefly discussed in this chapter, the pseudo-probabilistic non-exclusive anomaly classification capability provided by the MsCNN's softmax layer should be further explored.

Critically, the presented approach performs layer-wise anomaly classification at a speed fast enough to enable its implementation in a real-time environment; such an implementation is further discussed in Chapter 5. Finally, the author does not consider it possible to directly extend this work to the detection of micron-scale flaws, such as those discussed in Chapters 2

149

and 3, using a single, stationary camera⁸⁰ monitoring the entire powder bed. Indeed, it is this fundamental limitation which provided significant motivation for the melt pool monitoring work presented as the final topic of this thesis (Chapters 6 and 7).

⁸⁰ For a visible-light camera, the physical sensor pixel can be no smaller than the wavelength of visible light (0.4 μ m to 0.7 μ m [205]). As a result, assuming an ideal camera with 1 μ m square pixels, achieving a 10 μ m spatial resolution across the 250 mm × 250 mm EOS build plate would require a 625 MP sensor, at least 25 mm × 25 mm in size. Furthermore, if reliable detection of a flaw requires approximately 10 pixels × 10 pixels of data (a reasonable assumption based on the *patch* sizes used in this chapter) then the sensor size increases to that of the build plate itself.

5 Topic 3: L-PBF Build Analyses using Autonomous Powder Bed Anomaly Detection and Classification

5.1 Background

The development of the ML methodologies presented in Chapter 4 enabled the fruitful analysis of multiple L-PBF builds performed on the EOS M290 at CMU's NextManufacturing Center. This chapter presents analyses from ten case studies, each of which focuses on a unique aspect of part quality or build stability. Additionally, several of these analyses have been used by both internal and external users of the EOS M290 to understand build failure modes and redesign components for future builds. Particular attention is paid to the discussion of geometries which are traditionally difficult to manufacture additively, such as overhangs [182] and thin wall structures [183]. Correlations between fusion processing parameters (e.g. laser beam power, travel speed, and hatch spacing) and powder bed anomalies are studied. The impact of non-standard material systems on anomaly detections is explored and a mid-build malfunction of the EOS M290 is identified. Finally, potential strategies for real-time mitigation of powder bed anomalies are briefly discussed.

Importantly, many of the conclusions presented in this chapter could not have been reached without the autonomous analysis of in-situ process monitoring data. Many of the detected flaws occur within the volume of the parts and could not be identified ex-situ; furthermore, the data sets are far too large (many thousands of layers) for manual analysis of the powder bed images to be feasible. While the work in Chapter 4 was originally envisioned as a step toward real-time process feedback control, the developed methodology has also proven to be a powerful offline data analytics tool.

151

All of the analyses presented in this chapter were produced using the final Multi-scale Convolutional Neural Network (MsCNN) presented and characterized in Chapter 4. Many of the analysis visualization modalities in this chapter are presented without pretext; please refer to Section 4.6 for details regarding the interpretation of these figures. While the work presented in this chapter was not directly supported by any entity, the funding sources and collaborators for each of the presented case studies are indicated in the body of the text.

5.2 Results

5.2.1 Overhangs and Sudden Delamination

Many of the process stability issues observed during L-PBF are related to mid-build deformation of the printed parts due to residual thermal stresses [118]. Such stresses can result in the warping of an unsupported overhang region up above the powder bed as well as the sudden delamination of a part from its support structures or the build plate itself. This subsection presents two examples of each thermal stress-induced failure mode across three different case studies.

In April of 2016, several robotic arm components were built out of AlSi10Mg for Ben Brown of CMU's Robotics Institute. Figure 5.1 shows a Computer Aided Design (CAD) rendering of one of the components which possesses a relatively large unsupported overhang. Unfortunately, these overhanging regions began to warp and were impacted by the recoater blade – resulting in a failure of the entire build. Figure 5.2 shows several layers both before and after the warping began. Note that several of the components were manually deactivated by the operator after layer 520.

152



Figure 5.1: A CAD rendering of the robotic arm component. Eight of these components were built out of AlSi10Mg on a single build plate.



Figure 5.2: Images of the powder bed ranging from 15.21 mm to 15.69 mm above the build plate. Powder is first spread over the overhang region at layer 508. The three furthest-right components were manually deactivated by the operator after layer 520 in an attempt to preserve the remainder of the build. The entire build was manually aborted by the operator after layer 523.

In a real-time implementation of the MsCNN (see Section 5.3), the detection of a relatively low severity anomaly may be used to automatically trigger mitigation strategies in order to prevent the occurrence of more severe anomalies in subsequent layers. Toward that goal, Figure 5.3 demonstrates a correlation (in this build) between detections of *super-elevation* and detections of *part damage* in subsequent layers. Specifically, Figure 5.3 plots the percentage of *super-elevation* pixels at each layer which are later classified as *part damage*. For example, 66% of the pixels classified as *super-elevation* in layer 508 are classified as *part damage* in at least one layer between layers 509 and 523. Note that the downward trend in "future" *part damage* "predictions" is due to the fact that the five components on the left-hand side of the build plate did not fail before the build was manually aborted by the operator after layer 523 (see Figure 5.2). It is considered likely by the author that the remainder of the components would have eventually failed had the build been allowed to continue. Finally, note that the proper design of feedback control heuristics to enable mitigation of "future" *part damage* would require a substantial amount of additional research in this area (see Section 8.3).



Figure 5.3: A simplified *global build report* showing the detections of *super-elevation* and *part damage* between layer 500 and layer 523. The vertical part profile is shown in green (see Section 4.6.3 for details) and indicates the presence of an unsupported overhang at layer 508. The line labeled "Future Part Damage" reports the percentage of *super-elevation* pixels at each layer which are eventually classified as *part damage* in a subsequent layer.

In June of 2016, CMU's NextManufacturing Center was commissioned to build two models⁸¹ of CMU's Hamerschlag Hall out of Ti-6Al-4V (Ti64) for the groundbreaking of the new ANSYS Mechanical Engineering building. The build was performed using a chamber preheat of 35 °C and a (nominal) powder layer thickness of 30 µm. Figure 5.4 shows a CAD rendering of one of the models. This build provides an example of *part damage* due to relatively small scale overhangs as well as sudden delamination of a part from its support structures. Figure 5.5 shows a simplified *global build report*; anomaly detection spikes are visible at several layers; two of these layers, 760 and 1,954 are discussed in detail below.

⁸¹ The models were originally designed by Michelle Ma, Eric Myers, and Ryan Pearce for their final project in the course "Additive Manufacturing for Engineers" (24-210) at CMU and later modified by Dr. Zachary Francis, also of CMU.





The spike in *part damage* classifications at layer 760 (Figure 5.6) is a detection of a sudden delamination that occurred during the build. At this layer, the residual thermal stresses [118] within the right Hamerschlag Hall model were sufficient to tear it from some of the supports anchoring it to the build plate. This sudden delamination "tossed" powder off of the part, and the resultant severe lack of powder is correctly classified as *part damage* (magenta). The disturbances to the surrounding powder bed are primarily, and correctly, classified as *debris* (white). However some the *patches* are incorrectly classified as *incomplete spreading* which can

be visually similar to *debris*. Note that the left Hamerschlag Hall model does not experience delamination.

At layer 1,954 (Figure 5.7) the algorithm detects multiple instances of *part damage* (magenta) which correspond to observed collapses in the rotunda windows (overhang regions) in the final part. The recoater impacted the rotunda region leading to detectable levels of *recoater hopping* (vertical teal line). *Part damage* is primarily detected on the right Hamerschlag Hall model and this is consistent with the post-build analysis of the final parts. This difference is particularly observable in the *heat map* shown in Figure 5.8; note that the rotunda on the right model has a high percentage of instances of *part damage*. The other regions where a large number of instances of *part damage* are detected correspond to the open doorways of both models; these doorways represent significant unsupported overhangs and the quality of the final part was noticeably poor in those locations⁸².

⁸² The Hamerschlag Hall models were built for an external customer and are no longer available to the author. No appropriate images of the as-built models were taken at the time the models were built.





Figure 5.6: Layer 760 (22.80 mm above the build plate). The green pixels show the CAD geometry outline of the parts at that layer.

Figure 5.7: Layer 1,954 (58.62 mm above the build plate). The green pixels show the CAD geometry outline of the parts at that layer.



Figure 5.8: A cumulative *heat map* showing the percentage of layers (throughout the build height) in which *part damage* was detected at each pixel. The positions of the parts on the build plate are shown as white outlines of their footprints. The sporadic white dots are artifacts present in the EOSPRINT environment which were captured during extraction of the part geometries (Section 4.2.5).

In January of 2018, Brian Fisher of CMU conducted two builds of identical turbine blades made out of Ti-6AI-4V. One build was performed using the EOS nominal chamber preheat of 35 °C while the other build was performed at a chamber preheat of 500 °C; the higher preheat was made possible through the use of a custom "heated build plate insert" designed by Fisher [34]. These two builds were performed in order to demonstrate reduced residual stresses at higher background temperatures [34]. Ex-situ observations by Fisher determined that while the turbine blade with the weakest support structure delaminated during the 35 °C build, no such failure occurred for the corresponding turbine blade during the 500 °C build [34]. These exciting results are also borne out by the MsCNN analyses of the in-situ powder bed images. Figure 5.9 shows the *local build reports* for the relevant turbine blade during each build. While a sudden spike in anomaly detections due to a large disturbance of the powder bed (Figure 5.10) is visible at the time (layer 386) of delamination during the 35 °C build, no such spike is observed for the 500 °C build. While delamination did not occur at the higher chamber preheat temperature, a substantial increase in detections of *super-elevation* is evident. The author hypothesizes that these detections are the result of swelling [115], [147, p. 28] of the turbine blade. The relationship between swelling and *super-elevation* is discussed in detail in Section 5.2.4.



Figure 5.9: The delamination event during the low preheat temperature (35 °C) turbine blade build is indicated by the spike in anomaly detections at layer 386. Note that this delamination occurred well above the support-part interface at 2.55 mm of build height. Also note the increased number of *super-elevation* detections during the high preheat temperature (500 °C) turbine blade build. The detections of *incomplete spreading* and *debris* during the 500 °C turbine blade build are artifacts caused by the "heated build plate insert" described by Fisher [34].



Figure 5.10: Layer 386 (11.58 mm above the build plate) of the low preheat temperature (35 °C) turbine blade build. The turbine blade of interest lies within the blue bounding box. The delamination event resulted in a significant disturbance in the powder bed, similar to that observed in Figure 5.6. The green pixels show the CAD geometry outline of the parts at that layer.

It is evident from the part profile (green line) reported in Figure 5.9 that the delamination event occurred well above (i.e. at a later layer than) the support-part interface (11.58 mm versus 2.55 mm above the build plate). Note that this same behavior was observed during the delamination of the Hamerschlag Hall model discussed previously. Therefore, while the "timings" of the delamination events were properly captured by the MsCNN analysis, the layer at which structural failure occurred could not be identified and flagged by the algorithm. Indeed, this is a fundamental limitation of using the analysis of powder images to identify spatially-accurate macro-scale defect locations for part quality assurance applications. Furthermore, it is likely that the delamination event itself will become more difficult to detect via powder bed imaging the greater the depth of powder separating the delamination layer from the surface of the powder bed. In such cases, passive acoustic monitoring [130] of the build process may be an appropriate defect detection alternative.

5.2.2 High-Aspect Ratio Components

In September 2016, multiple "dog bone" tensile specimens⁸³ were built out of stainless steel 316L (SS 316L), in varying orientations and with different support schemes, for the Bettis Naval Nuclear Laboratory. The build was performed using a chamber preheat of 80 °C, a and (nominal) powder layer thickness of 20 µm. Figure 5.11 shows a 3D reconstruction of the entire build. In this subsection, two different tensile specimens (one oriented horizontally, the other vertically) are discussed in detail. Note that while all 25 of the tensile specimens and cylindrical witness coupons are visible as green outlines in Figures 5.13 and 5.16, the specific tensile specimens being discussed are bounded by blue boxes.

⁸³ The tensile specimens were designed by Dr. Colt Montgomery of CMU.



Figure 5.11: A 3D rendering of the tensile specimen build with the detected *part damage* highlighted in magenta. The horizontally-oriented tensile specimen used minimal support structures underneath

the gage section in an attempt to reduce their influence on measured part properties during mechanical testing. From the *local build report* in Figure 5.12, it is clear that significant *part damage* (magenta, Figure 5.13) is detected as soon as the component transitions from the support material to the tensile specimen itself. This occurs at 5.06 mm of build height, or layer 253. The lower subplot (green) of Figure 5.12 shows the percentage of pixels within the blue bounding box (Figure 5.13) that lie on top of the tensile specimen itself as a function of build height. Recall that an increase in this percentage represents an overhanging region, and a sharp increase implies a largely unsupported overhang. It is inferred that the support structures were insufficient for the substantial overhangs presented by the horizontally-oriented tensile specimen. As a result, failure of the supports occurred once the recoater blade passed back over that region of the powder bed, striking the first layer of the tensile specimen itself (i.e. above the supports). Figure 5.14 shows the as-built horizontally-oriented tensile specimen.



Figure 5.12: Local build report and vertical part profile for the horizontally-oriented tensile specimen of interest.



Figure 5.13: (left) Note the instances of *part damage* (magenta) detected at layer 253 (5.06 mm of build height) within the blue bounding box. The green pixels show the CAD geometry outline of the parts at that layer. Figure 5.14: (right) The as-built horizontal tensile specimen with visible defects highlighted by a red bounding box. Note that the regions encompassed by the blue and red bounding boxes are nominally identical.

From Figure 5.15 it is apparent that the vertically-oriented tensile specimen built successfully throughout most of its height, but instances of *part damage* (magenta) are suddenly detected around layer 3,524 (70.48 mm of build height). This corresponds to the tensile specimen transitioning from the gage section into the upper grip section, a period during which a large overhang is produced. Based on Figures 5.16 – 5.18, it is surmised that the unsupported overhang warped upwards (due to residual thermal stress [118]) enough to be impacted by the recoater blade. This impact is highlighted by the extensive *recoater hopping* classifications (vertical teal line). The impact caused the tensile specimen to bend and then "spring back," tossing powder away and leaving a powder cavity to the right of the tensile specimen (highlighted by *debris* classifications in white). This cavity prevented proper powder coverage of this region in subsequent layers, worsening the situation. Analysis results, including those presented in this subsection, informed many of the design changes incorporated into a subsequent build at CMU's NextManufacturing Center, which was also supported by the Bettis Naval Nuclear Laboratory.



Figure 5.15: *Local build report* and vertical part profile for a vertically-oriented tensile bar.



Figure 5.16: (left) Note the instances of *part damage* (magenta) detected at layer 3,524 within the blue bounding box. Also note the *debris* (white) classifications indicating the powder cavity to the right of the vertical tensile specimen. The green pixels show the CAD geometry outline of the parts at that layer.

Figure 5.17: (center) A 3D rendering of the vertical tensile specimen with *part damage* highlighted in magenta. Figure 5.18: (right) The as-built vertical tensile specimen with the visible defects bounded by a red box.

5.2.3 Critical Orientations of Thin Wall Structures

In October 2016, an impeller⁸⁴ was built out of stainless steel 316L (SS 316L) to demonstrate the capabilities of the EOS M290 and as part of CMU's NextManufacturing Center industry training initiative. The build was performed using a chamber preheat of 80 °C and a (nominal) powder layer thickness of 20 μ m. The thin impeller blades and thin powder layers make this part a challenge to build successfully. Figure 5.19 shows the as-built impeller blade. Note that while one half of the impeller built correctly (Type A blade orientation), the blades on the other half collapsed (Type B blade orientation). The MsCNN was employed to determine the possible reason(s) for the partial build failure and identify any potential strategies for improving build quality in the future.



Figure 5.19: The as-built impeller after being cut in half with a Wire Electrical Discharge Machine (EDM) to preserve only the half that built correctly.

⁸⁴ The original CAD model of the impellor was downloaded from an online source by Dr. Colt Montgomery of CMU.

While it was immediately clear that the thin impeller blades failed due to repeated impacts with the recoater blade, it was initially unclear why only the impeller blades on one half to fail. Figure 5.20 indicates that *super-elevation* occurred on both halves of the impeller – suggesting that the recoater blade likely impacted most of the impeller blades, including many that did not ultimately fail. Figure 5.21 highlights the impeller blades that collapsed during the build, note that only impeller blades with their leading-edge pointed away from the direction of the incoming recoater blade (Type B) failed to build correctly.



Figure 5.20: A cumulative *heat map* showing the percentage of layers (throughout the build height) in which *super-elevation* was detected at each pixel. The positions of the parts on the build plate are shown as white outlines of their footprints.

Figure 5.21: A cumulative *heat map* showing the percentage of layers (throughout the build height) in which instances of *part damage* were detected at each pixel. The positions of the parts on the build plate are shown as white outlines of their footprints.

Based on Figures 5.20 and 5.21, the author hypothesizes that some plastic deformation occurred when the impeller blades were struck by the recoater blade. For the impeller blades that built correctly (Type A), this deformation would be in approximately the same direction as the shift in blade CAD geometry in the subsequent layer, while the deformation and geometry shift would be in opposite directions for Type B blades. In this context, "geometry shifts" refer to the discrete, layer-wise, in-plane shifts that are used to construct the smooth 3D curvature (defined by the CAD geometry) of the impeller. As a result, plastic deformation inflicted on the Type A impeller blades would be less likely to cause cascading failures in subsequent layers than the same deformation inflicted on the Type B impeller blades. Figure 5.22 shows this hypothesis schematically.



Figure 5.22: A schematic representation of the two extreme impeller blade orientations. In the case of the Type A blade orientation, any plastic deformation due to a recoater strike occurs in the same direction as the geometry shift in the subsequent layer. In the case of the Type B blade orientation, the opposite is true – plastic deformation occurs in the opposite direction as the geometry shift in the subsequent layer. As a result, the Type B impeller blades are less likely to recover from a recoater strike, as the following layer will likely be largely deposited on a bed of unfused powder, with no substantial connection to the previous layer of the impeller.

Because of this powder bed anomaly analysis, only half-impellers (of Type A) were built for the subsequent industry training event; complete collapses of the impeller blades due to recoater impacts were not observed for this subsequent build. Successfully building an entire impeller of this design would likely require a combination of approaches which might include: increasing the thickness of the impeller blades, increasing the powder layer thickness (to reduce the likelihood of recoater impacts due to *super-elevation*), and modifying the process parameters (e.g. laser beam power and laser beam travel velocity) to reduce swelling of the part [147].

5.2.4 Impact of Fusion Processing Parameters on Powder Layer Deposition

Much of this chapter focuses on the effects of part geometry and orientation on the appearance of the powder layer. Intriguingly, the parameters used to fuse the part geometry can also impact the powder spreading process. For example, a correlation between the laser scan strategy and the spatial distribution of spatter particles is identified in Section 3.3.2. In this subsection, preliminary correlations between the occurrence of *super-elevation* and laser beam power, travel velocity, and hatch spacing are observed for two builds.

In August of 2017, members of CMU's NextManufacturing Consortium (including General Electric, General Motors, and Arconic) supported a build of 40 tensile specimens⁸⁵ (Figure 5.23) in the EOS standard AlSi10Mg material system. A total of eight tensile specimens were built using each of the five process parameters combinations⁸⁶ enumerated in Table 5.1. The build was performed with a chamber preheat of 200 °C, a (nominal) powder layer thickness of 30 µm, and a (nominal) beam diameter of 100 µm. From Figure 5.24 it is immediately evident that detections of *super-elevation* were far more prevalent for some tensile specimens than others. Furthermore, the detected instances of *super-elevation* are indeed representative of the true situation; as shown in Figure 5.25 the edges of the tensile specimens are visibly extending above spread powder layer.

⁸⁵ The tensile specimens were designed by the author and Brian Fisher of CMU.

⁸⁶ The process parameters were chosen by Brian Fisher and Dr. Sneha Prabha Narra of CMU.



Figure 5.23: A CAD rendering of a tensile specimen with a non-contact thermal support separated from the specimen by 500 μ m throughout the height of the build. Forty identical tensile specimens were built out of AlSi10Mg on a single build plate.

Table 5.1: Process parameter combinations for the AlSi10Mg tensile specimens.

Parameter Number	Beam Power (W)	Beam Velocity (mm/s)	Hatch Spacing (µm)	Layer-wise Energy Density (mJ/mm ²)
1	370	1300	190	1.50
2	370	1300	220	1.29
3	370	600	230	2.68
4	250	1000	100	2.50
5	370	1300	110	2.59



Figure 5.24: A cumulative *heat map* showing the percentage of layers (throughout the build height) in which *super-elevation* was detected at each pixel. The positions of the parts on the build plate are shown as white outlines of their footprints.

Figure 5.25: Layer 3,483 (104.49 mm above the build plate). The green pixels show the CAD geometry outline of the parts at that layer. Note that the circular perimeters of some tensile specimens are visible above the powder layer.

When plotted in a process space defined by only the beam power and velocity, no trends in the *super-elevation* detections are apparent. However, when the average detections (throughout the height of the build) of *super-elevation* are plotted with respect to the layerwise energy density⁸⁷ (5.1), as shown in Figure 5.26, a correlation is apparent. This correlation can be quantified using a Pearson rank correlation test [108]. The correlation coefficients⁸⁸ (p) for detections of *super-elevation* and *part damage* with respect to layer-wise energy density are 0.93 and 0.70, respectively. The corresponding p-values are 0.021 and 0.18, respectively. Because *super-elevation* can often be associate with, and miss-classified as, *part damage* (Sections 4.2.4 and 4.5.3) it is worthwhile to also consider the correlation coefficient and p-

⁸⁷ This avenue of exploration was initially suggested by Brian Fisher of CMU.

⁸⁸ A Pearson rank correlation test quantifies how well the relationship between two variables follows a monotonic function. In this application of the correlation test, high correlation coefficients and low p-values indicate strong correlations between the particle size characteristics and the quality metrics. The correlation coefficient ranges from $\mathbb{R}[-1, 1]$, with a value of 1 indicating a monotonically increasing relationship and a value of -1 indicating a monotonically decreasing (i.e. inverse) relationship [108]. Note that the reported correlation metrics are for the mean measurement values.

value for the combination of the two anomaly detections which are 0.94 and 0.016, respectively.

$$u = \frac{P}{vh}$$
(5.1)

Where u is the layer-wise energy density, P is the beam power, v is the beam travel velocity, and h is the hatch spacing.



Figure 5.26: Average detections of *super-elevation* and *part damage* throughout the build height as a function of layer-wise energy density. The error bars represent 95% confidence intervals about the mean value; where each mean value is an average of the anomaly detections for the eight tensile specimens built with each process parameter combination. The three sets of tensile specimens built with the EOS nominal beam power and travel velocity (but various hatch spacings) are indicated by the cross marks. The EOS nominal hatch spacing is 190 µm.

Similar behavior to that described above was also observed during a build performed in October 2016 by Dr. Sneha Prabha Narra of CMU. In this case, the build consisted of 30 blocks (similar to the MLP samples introduced in Section 2.2.1) of material built with the EOS standard Inconel 718 (In718) material system. Each block of material was exposed using a different set of process parameters (laser beam power, travel velocity, and hatch spacing) as reported in [12, Fig. 6.7]. The build was performed using a chamber preheat of 80 °C, a (nominal) powder layer thickness of 20 µm, and a (nominal) beam diameter of 100 µm. From Figure 5.27 it is

immediately evident that detections of *super-elevation* and *part damage* were far more prevalent for some blocks than others. Furthermore, the detected instances of *super-elevation* and *part damage* are indeed representative of the true situation; as shown in Figure 5.28 the corners of the blocks are visibly extending above spread powder layer. Note that due to significant confusion between detections of *super-elevation* and *part damage* for this build (the *super-elevated* corners are visually similar to small-scale *part damage* defects, see Sections 4.2.4 and 4.5.3) the following analyses are based on the combined detections of the two anomaly types.





Figure 5.27: A cumulative *heat map* showing the percentage of layers (throughout the build height) in which *super-elevation* or *part damage* was detected at each pixel. The positions of the parts on the build plate are shown as white outlines of their footprints.

Figure 5.28: Layer 223 (4.46 mm above the build plate). The green pixels show the CAD geometry outline of the parts at that layer. The outermost square part outline is an artifact present in the EOSPRINT environment which was captured during extraction of the part geometries (Section 4.2.5).

When plotted in a process space defined by only the beam power and velocity, no trends in the *super-elevation* or *part damage* detections are apparent. However, when the average detections (throughout the height of the build) of *super-elevation* and *part damage* are plotted with respect to the layer-wise energy density, as shown in Figure 5.29, a weak correlation may be present. As quantified by a Pearson rank correlation test, the correlation coefficient and pvalue are respectively 0.46 and 0.010 for detections of *super-elevation* and *part damage* with respect to layer-wise energy density.



Figure 5.29: Average detections of *super-elevation* and *part damage* throughout the build height as a function of layer-wise energy density. Each data point represents one of the 30 sample blocks. Note that the EOS nominal process parameters for In718 using 20 μ m powder layers (and the corresponding layer-wise energy density) are not available to the author.

In the case of the AlSi10Mg tensile bars there is a strong and statistically-significant correlation between the layer-wise energy density and detections of *super-elevation*. This statistical relationship is further bolstered by the fact that three of the five studied parameter combinations have the same beam power and travel velocities and only differ with regards to their hatch spacing. This suggests that it is indeed the layer-wise energy input which is affecting the surface of the fused parts and not just the beam power and travel velocity. A substantially weaker correlation is observed between powder bed anomalies and layer-wise energy density in the case of the In718 blocks; however note that while weak, the correlation is still statistically-significant. Notably, the link between energy density and non-overhang part deformation has been explored by Sames [147, pp. 160–167] who referred to the flaw as
"swelling." Swelling was observed ex-situ for Inconel 625 parts fused using high-energy densities in the Arcam EB-PBF process [147, Fig. 77]. It is believed to be related to surface tension effects within the melt pool (as well as the surrounding area) when the local background temperature is elevated [115], [147, p. 28].

Indeed, this correlation, and swelling defects in general, should be studied further as there is limited prior work in the literature. Furthermore, swelling severe enough to lead to *super-elevation* can indeed become a concern for process stability, as indicated by the recoater strikes (*recoater hopping*) visible in Figure 5.25. Interestingly, swelling during the AlSi10Mg tensile specimen build appears to be limited to the side of the specimens furthest from the non-contact support structure (Figures 5.23 and 5.24). As those structures were intended to draw heat away from the specimens, it is possible that their presence reduced the local background temperatures enough to mitigate swelling even during exposure using high energy density process parameters.

Finally, a periodicity in the combined *super-elevation* and *part damage* detections was observed throughout the build height of the In718 blocks. A Fourier frequency analysis yields strong peaks at 3.90 layers/anomaly-peak and 2.69 layers/anomaly-peak as shown in Figure 5.30. As discussed in Section 3.3.2, such periodicities may be related to the default EOS M290 laser scan strategy (Figure 1.8) rotation of 67° every layer [41]. Operating under this assumption, these layer-wise periods translate to respective scan strategy rotation periods of 260°/anomaly-peak and 180°/anomaly-peak. However the presence of two peaks in frequency space suggests that the critical orientation of the scan strategy may precess during the build.



Figure 5.30: The results of a Fourier frequency analysis of the combined detections of *super-elevation* and *part damage* throughout the height of the build containing the In718 blocks. Two peaks in frequency space are prominent.

This behavior is substantially more complex than that observed in Section 3.3.2 and its detailed study is beyond the scope of this work. Nonetheless, the author speculates that the observed periodicity may be the result of an interaction between the stripes and the corners of the blocks. Figure 5.31 demonstrates that at certain orientations (Case A) the geometry of the block does not influence the stripe width whereas in other orientations (Case B) the stripe width narrows as it approaches the corners. Such a narrowing of the stripe width is expected to increase the local background temperature [184] and increase the potential for swelling to occur. Indeed, this corner behavior (for a non-rotating scan strategy) was also observed ex-situ by Sames [147, Fig. 77]. The effects stripes intersecting with part geometry are explored further in Chapter 7.



Figure 5.31: Examples of two different interactions between the laser scan strategy (stripes) and the corners of the In718 blocks. The Case B situation (approximately) repeats every four layers. Observe that a similar situation will also be present when each of the Case B layers has rotated approximately 180°.

5.2.5 Material Systems with Distinct Appearances

The MsCNN algorithm has been used to analyze several builds making use of non-standard material systems. One example of this is reported in Section 3.3.2 and another example is discussed, briefly, below. Bronze is not currently a material system supported by EOS [27]; in November of 2016, Matthews International Corporation (Bronze Division) supported two builds on the EOS M290 at CMU's NextManufacturing Center using non-standard bronze powder. Despite the dramatically different visual appearance of bronze powder compared to other metal powders (Figure 5.32), the algorithm performed robustly, with no re-training required. This robustness may be partially the result of the substantially reduced visual difference between bronze powder and other metal powders when viewed in grayscale (Figure 5.33). Figure 5.34 shows a 3D rendering of the anomalies detected during the construction of a model DNA helix, just 10 mm in diameter. Unfortunately, the relatively fragile helix failed to build correctly (due to impacts by the recoater blade), but this failure was successfully identified by the MsCNN algorithm.





25 50 75 100 125 150 175 200 225 25 X (mm) ← recoater direction

Figure 5.32: A color image taken of the bronze powder after completion of the first build supported by Matthews International Corporation. Note the characteristic yellow/gold color. Figure 5.33: A grayscale image captured by the EOS M290's powder bed camera during the second build supported by Matthews International Corporation. This image was taken at layer 685 (20.5 mm of build height) and also shows the MsCNN anomaly classifications. The green pixels show the CAD geometry outline of the parts at that layer.

Part
 Outline

Recoater

Hopping Recoater

Streaking

Spreading

Debris

Super-Elevation

Part Damage Incomplete



Figure 5.34: A 3D rendering of the model DNA helix manufactured during the second build supported by Matthews International Corporation. Instances of *part damage* are highlighted in magenta.

While the trained MsCNN algorithm has proven itself robust during analyses of builds using a wide range of material systems, several behavioral differences between material systems have been observed. Most notably, powder particle size has been correlated with powder bed anomaly detections as described in Section 3.3.2. Also of interest, the EOS standard⁸⁹ AlSi10Mg material system has a higher level of background (i.e. not related to the parts being built) anomalies than the other standard material systems as reported in Table 5.2. The reported layer-wise anomaly detection averages are based on rectangular regions of the powder bed ranging in size from 2500 mm² to 6250 mm² and extending from the extreme upper-right corner⁹⁰ of the build area. Each region was selected such that part fusion would not be expected to influence the appearance of the powder bed within the region; Figure 5.35 shows an example selected region. Finally, full powder layer coverage of the build plate requires between 22 layers and 61 layers for the AlSi10Mg material system while less than five layers are required for the other material systems studied. Figure 5.36 shows the simplified *global build reports* for the four representative AlSi10Mg builds studied in this subsection.

⁸⁹ Note that all of the AlSi10Mg builds studied in this chapter were built using EOS standard AlSi10Mg powder purchased prior to July 31st, 2017. The AlSi10Mg powder currently (c.a. 2018) available through EOS has not been studied by the author.

⁹⁰ The upper-right corner of the build area is expected to be the region of the powder bed least affected by part fusion. As the argon shielding gas flow is antiparallel to the *y*-axis, spatter is less likely to fall along the upper edge than elsewhere on the powder bed. Because the recoater blade spreads powder in the direction antiparallel to the *x*-axis, debris from damaged parts is less likely to be deposited along the right edge than elsewhere on the powder bed. See Figure 1.4 for additional clarification.

Material System	Combined Average Debris and Recoater Streaking Detections	Number of Builds Analyzed
EOS AlSi10Mg	1.3% - 2.8%	4
EOS Stainless Steel 316L	<0.05% - 0.1%	4
EOS Ti-6Al-4V	<0.05%	2
EOS Inconel 718	<0.05%	4
EOS Inconel 625	<0.05%	1
Matthews International Corporation Bronze	<0.05%	2

Table 5.2: Background powder bed anomaly detections for selected material systems.



Figure 5.35: The selected region of the powder bed is shown in light gray and extends from 225 mm to 250 mm along the *x*-axis and from 0 mm to -250 mm along the *y*-axis.



Figure 5.36: *Global build reports* from four different AlSi10Mg builds showing the detections of *recoater streaking* and *debris*. The layers at which full powder coverage of the build plate is achieved are indicated by solid vertical black lines. Full powder coverage is considered to occur once the anomaly detections have reached their (approximatly) steady state value based on visual observation. Note that for clarity, the *y*-axes of each subplot do not span the same detection range.

5.2.6 Detection of L-PBF Machine Malfunctions

In a real-time monitoring implementation, the MsCNN algorithm has the potential to detect certain malfunctions of the L-PBF machine. In July of 2016 a heat exchanger (similar to that presented in Section 4.6.1) was built for Dr. Samikshya Subedi and Prof. Anthony Rollett of CMU's Materials Science and Engineering Department under a US Department of Energy grant (DE-FE0024064) and in collaboration with Oregon State University and Prof. Vinod Narayanan at the University of California Davis. Unfortunately, at layer 619 the EOS M290 machine failed to raise the powder dispenser before spreading the next layer. The root cause of this failure remains unknown although a similar issue occurred during the PS #5 as discussed in Section 3.3.1. As no powder was spread over the parts, a dramatic increase in *super-elevation*

detections is evident in Figures 5.37 and 5.38. While the EOS M290 began spreading powder again starting at layer 655 the initial disruption eventually lead to a complete failure of the build at layer 668 as shown in Figure 5.39.



Figure 5.37: A *global build report* showing detections of *super-elevation* and *part damage* throughout the build. The powder spreading failure is clearly visible as a rapid increase in the detections of *super-elevation*.



Figure 5.38: Layer 619 (12.38 mm above the build plate). Severe *super-elevation* is indicated in red. The green pixels show the CAD geometry outline of the parts at that layer.



Figure 5.39: Layer 668 (13.36 mm above the build plate). Severe *part damage* is indicated in magenta. At this layer, the powder bed is approximately 1 mm lower than the powder dispenser (versus the nominal gap of 20 μ m). As a result, insufficient powder is fetched from the powder dispenser to completely cover the build area. Indeed, the maximum horizontal (*x*-axis) extent of the powder spread is visible as the vertical line of *debris* classifications in white.

5.3 Proposed Anomaly Mitigation Strategies for Real-Time Implementations

As observed in Section 4.5.4, the final MsCNN methodology is capable of classifying an entire powder bed image in approximately the same amount of time it takes for the EOS M290 to spread a single layer of powder. Therefore there are no fundamental impediments to implementing powder bed anomaly classification on-line, in real-time. Unfortunately, restrictions imposed by the machine manufacturer prevent trivial real-time acquisition of images from the powder bed camera and completely prevent automatic changes to the EOS M290's operation during a build. Such restrictions are not imposed by all L-PBF machine manufactures with the same degree of strictness (and EOS GmbH may reduce their own restrictions in the future), therefore it is worthwhile to consider a hypothetical real-time implementation of the methodology presented in Chapter 4.

The logical "first step" toward a real-time implementation would be utilizing anomaly classification in a purely supervisory capacity. For example, an operator could be automatically notified if anomaly detections in a single layer (or over the course of multiple layers) exceeded a specified threshold. It is even conceivable that, given an appropriate dataset, ML could be employed to identify the appropriate notification thresholds based on historical operator behavior. In other words, operators manufacturing aerospace components may desire a lower notification threshold than those manufacturing non-safety-critical components. Beyond monitoring the powder bed, there are also several potentially viable methods for mitigating flaws detected during the build.

Incomplete spreading is trivially addressed by increasing the amount of powder spread each layer (the dosing factor). Indeed, ConceptLaser's QM Coating system is now on the market and automatically corrects for *incomplete spreading* [146]. Similarly, isolated instances of *recoater streaking* can often be mitigated by re-spreading the powder layer. If, however, *recoater streaking* is detected in the same region of the powder bed over the course of several layers it may be an indication that the recoater blade should be cleaned or replaced. While cleaning the recoater blade currently requires operator intervention in all L-PBF machines known to the author, it could hypothetically be performed robotically with relative ease.

If the detected instances of *part damage* for a given part exceed a specified threshold (beyond which the part designers, process engineers, and operators believe the part will not be functional and/or safe to use) that part could be automatically "deactivated," allowing the rest of the build to continue without wasting resources (e.g. build time and feedstock) on a part destined to fail (Section 5.2.1). Similarly, deactivating a failing part may prevent *debris* from impacting "downstream" parts which are otherwise undamaged. Figures 5.13 and 5.16 show examples of *debris*, originating from a damaged part, being dragged from right-to-left across the powder bed by the recoater blade during powder spreading. This *debris* behavior is particularly visible when viewing layer-by-layer time-lapses of the building process (Section 4.6.7).

If *super-elevation* or *recoater hopping* can be detected in real-time, it may be possible to avoid severe part-recoater blade impacts. In some situations, increasing the powder layer thickness for subsequent layers (and modifying the fusion processing parameters as in Chapter 3) may be viable. Adjusting fusion process parameters such as the stripe width (Figure 1.8) may

increase the local pre-heat temperature [184] and decrease the internal residual stresses [34], [118], [140] thereby mitigating *super-elevation* due to warping directly. Conversely, if the detected *super-elevation* is the result of swelling instead of warping, a reduction in layer-wise energy density (possibly by increasing the hatch spacing) may be appropriate [147]. Note that a substantial amount of research on this type of process control would be required to determine the viability of this mitigation strategy.

5.4 Discussion and Summary

In this chapter, the Multi-scale Convolutional Neural Network presented in Chapter 4 is used as a powerful data analytics tool to study the layer-wise powder spreading and fusion processes in an L-PBF machine. The results presented in this chapter not only provide a qualitative representation of the performance of the final MsCNN methodology, but also include several unique insights based on the autonomous analysis of thousands of powder layers. A total of ten case studies are explored in detail in order to draw conclusions relating to build geometry, fusion process parameters, laser beam scan strategy, powder spreadability, and process stability.

In Section 5.2.1, large unsupported overhangs are observed to warp upwards out of the powder layer due to residual thermal stresses. This deformation, originally classified as *super-elevation* is classified as *part damage* in subsequent layers after the warped regions of the parts are impacted by the recoater blade. Depending on the size of the overhangs and the severity of the warping, such deformation may result in the failure of the entire build. Delamination events were observed at the support-part interface during two builds. Delamination is also the result

of residual thermal stresses; however the event tends to occur during the fusion of a layer well above the layer at which the failure physically occurs – making detection a challenge. Identical parts (and support structures) were built at two different chamber preheat temperatures; while delamination was observed during the low temperature build it did not occur during the high temperature build. This difference in behavior can be attributed to the reduced level of residual thermal stress present within a part when fusion occurs at a higher background temperature [34].

In Sections 5.2.2 and 5.2.3, high-aspect ratio and thin wall part geometries are shown to have unique failure modes. In the case of thin, vertical tensile specimens, an impact from a recoater blade can elastically deform the part such that it "springs back" after the recoater blade passes over it. This motion can "toss" powder away from the tensile specimen and create a cavity which prevents proper powder coverage of the specimen in subsequent layers. The build stability of certain thin wall structures may be dramatically dependent upon their orientations relative to the motion of the recoater blade. Specifically, minor plastic deformation caused by an impact with the recoater blade may have either a minimal or a significant effect on the fusion and spreading of subsequent layers depending on the orientation of the part.

In Section 5.2.4, fusion parameters (laser beam power, travel velocity, and hatch spacing) are shown to correlate with detections of *super-elevation*. Specifically, a high layer-wise energy density results in increased detections of *super-elevation* due to swelling of the fused material. These observations are compared to ex-situ swelling observed in the literature [147] for the EB-PBF process. Of great interest, a correlation is also observed between swelling (quantified by detections of *super-elevation*), laser scan strategy, and part geometry. Specifically, the raster

length of the laser beam is dependent on part geometry (e.g. corners) only at certain stripe rotation angles and therefore only at certain layers. During fusion of such layers, the stripe width decreases near part corners thereby increasing the local background temperature and leading to an increased amount of detected part swelling.

In Section 5.2.5, the MsCNN is shown to perform robustly for a non-standard material system (bronze) with a significantly different visual appearance. The higher background rates of *recoater streaking* and *debris* detections for the EOS standard AlSi10Mg material system are also quantified. In Section 5.2.6 the MsCNN demonstrates the ability to detect a mid-build malfunction of the EOS M290 L-PBF machine. Finally, potential strategies for real-time mitigation of certain powder bed anomalies are proposed and discussed in Section 5.3.

6 Topic 4: Ex-Situ Melt Pool Morphology across Inconel 718 L-PBF Process Space

6.1 Background and Literature Review

Meaningful interpretation of data from in-situ process monitoring schema often requires an assessment of, and correlation with, data from ex-situ analyses. This is especially true for Machine Learning-based approaches including the in-situ melt morphology classification work presented in Chapter 7. This chapter describes the creation of a well-controlled ex-situ dataset designed specifically to enable the linkage of in-situ melt pool morphologies to process outcomes. In particular, the morphologies of interest in this work are those characteristic of process outcomes such as porosity formed by the keyholing mechanism [74], under-melting [74], and the surface tension-related balling phenomenon [87], [88]. Keyhole-mode melting occurs in the high energy density (high beam power, low beam velocity) region of process parameter space, where periodic vaporization of the molten material can occur. Under certain conditions, the resultant vapor pocket may become trapped as porosity in the solidified melt pool [74]. Under-melting occurs when the melt pool does not fully penetrate the powder layer [74]. Balling occurs in the higher beam velocity and beam power regime of process space [31]. Specifically, as the melt pools lengthen relative to their widths, Rayleigh instabilities driven by surface tension forces⁹¹ cause the tail of the melt pool to first form "humps" and eventually cease to be a continuous melt track, instead breaking up into discrete "balls" [88].

⁹¹ This phenomenon is also visible as a stream of water falls from spigot. The initially-contiguous water stream lengthens, and correspondingly narrows, as its velocity increases. Eventually, the diameter of the water stream becomes small enough relative to its length that surface tension forces are sufficiently high to break the continuous stream into discrete droplets of water.

This work constructs a statistically significant database of melt pool geometry and morphology across L-PBF process space for the EOS Inconel 718 (In718) alloy (Ni 50 – 55 wt%, Cr 17.0 – 21.0 wt%, Nb 4.75 – 5.5 wt%, Mo 2.8 – 3.3 wt%, Ti 0.65 – 1.15 wt% [102]). In718 is a nickel super-alloy with excellent high temperature performance and corrosion resistance [185], [186]. As a result, it has multiple applications in the aerospace [187] and energy sectors [186] but is difficult to machine via traditional processes [188]. Also of importance, prior work by Brian Fisher at CMU [34] confirmed that L-PBF In718 melt pools emit a sufficient amount of visible-spectrum light for the available high speed camera setup (Section 7.2.2) to collect useful data over a relatively wide range of camera settings.

Process mapping is a technique developed by Beuth et al. [32] that enables the correlation of process parameters to process outcomes (e.g. melt pool geometry, porosity, and as-built microstructure) by plotting those outcomes across process space. While process mapping [32] of L-PBF-processed In718 has already been performed by Dr. Sneha Prabha Narra of CMU [12, Ch. 6], several factors preclude the usage of those data for this application. One of the goals of the work presented in Chapter 7 is the observation of melt pools during the construction of "real" parts, a corollary of which is the observation of melt pools on top of a powder layer. Because melt pool morphology can vary significantly depending the presence of powder (vs. no powder) [36], [69] the use of the no-added powder experiments performed by Narra [12, Ch. 6] in this work would not be appropriate. One particularly relevant example of this powder dependence is the onset of the balling phenomenon which can occur over a greater range of processing parameters when powder is present [69]. Furthermore, the characteristic keyholing and balling morphologies are known to occur with varying degrees of periodicity within a single melt track [31], [85], [189]. This variability necessitates the analysis of a statistically-significant number of melt pool cross-sections if the ex-situ data are to be linked to the in-situ data. Such data are not available from the work of Narra [12, Ch. 6] as the melt pool geometry analyses are based on single cross-sections of each process parameter combination. Of additional interest, the combination of this database with the AlSi10Mg results from Chapter 2 allows for an in-depth discussion of the statistical distribution and variability of L-PBF melt pool geometry. While Francis [31, Fig. 4.5] and others [73] have characterized the variable behavior of melt pool depth in keyhole-mode melting, the existing literature is relatively sparse with regards to the statistical study of general melt pool variability in the AM processes. The work presented in this chapter was supported by CMU's Manufacturing Futures Initiative (MFI) (internal grant number 062900.005.105.100020.01).

6.2 Experimental Design and Methods

6.2.1 Build Conditions

The experiment described in this section is referred to as "Experiment 1 (E1)" in Chapter 7 to maintain a consistent nomenclature. The experiment was performed on an EOS M290 L-PBF machine at CMU's NextManufacturing Center and consisted of 10 single melt tracks at each of 36 different process parameter combinations (Table 6.1). The melt tracks were exposed on a 6 in \times 6 in \times 0.25 in In718 plate sourced from McMaster-Carr (P/N: 1099N8). Each melt track is

20 mm long to ensure steady-state⁹² melt pool conditions in the cross-sectioned regions. The layout of the single melt tracks is shown in Figure 6.1 and is such that all of the melt tracks for three parameter sets fit within a single sample puck (Section 6.2.2). Unlike the OLP experiments described in Chapter 2, the melt tracks in this chapter were exposed on top of a single layer of powder as opposed to a bare substrate and were spaced sufficiently far apart (500 μ m) to avoid any overlap between adjacent melt pools. Due to the importance of a consistent powder layer for these experiments, denudation of the region around the melt track is also an important consideration. Work by Matthews et al. [39] finds that the denudation zone extends between 100 µm and 200 µm from the center of the melt track for processing conditions similar to those explored in this chapter. Therefore the denudation zones of adjacent 1LSB melt tracks are not expected to have influenced the observed melt pool morphologies although a larger safety margin (spacing between adjacent 1LSB melt tracks) may be appropriate for future work. As such, they are designated 1LSB samples⁹³. To mitigate the effects of residual heating due to adjacent melt tracks [122] the exposure order of the 1LSB melt tracks was adjusted such that a minimum of 1.6 seconds elapsed between the lasing of adjacent melt tracks. Because the melt tracks are spatially distributed across the build plate, the high speed camera was not able to observe the lasing in-situ. The experiments used to collect the in-situ data are described in Section 7.2.

 $^{^{92}}$ Work by Fox [30] indicates that the response distance is proportional to the initial and final steady-state melt pool sizes. Specifically, it was found a melt pool should reach steady-state conditions within a travel distance of approximately three melt pool (final) depths. The deepest measured melt pools were on the order 500 μ m, suggesting a response distance of approximately 1.5 mm. Alternatively, the largest steady-state melt pool length predicted by the Rosenthal model (6.1) is on the order of 900 μ m. Both of these measures for response distance are substantially shorter than the approximately 10 mm distance from the end of the melt track to the crosssectioning point.

⁹³ One Layer Single Bead (1LSB) experiments, i.e. one layer of powder, single bead exposures.



Figure 6.1: The layout of the 1LSB melt tracks; there are 36 different beam power and velocity combinations. Note that the exposure order follows the parameter naming convention, that is, the first melt track of PV #1 was exposed first, followed by the first melt track of PV #2 through PV #36. After the first 36 melt tracks are exposed, the exposure sequence is repeated for the second through tenth melt tracks of each parameter combination. The four circles indicate the thru-holes used to mount the In718 plate to a modified steel build plate originally sourced from EOS.

All of the 1LSB melt tracks were exposed with a preheat temperature of 80 °C and a nominal⁹⁴ beam diameter of 100 μ m. The 36 processing parameter combinations (Table 6.1) were chosen based on prior L-PBF In718 work performed by Narra [12, Ch. 6] as well as an analytical heat transfer model of a moving point heat source known as the Rosenthal Equation [190]. The Rosenthal model (6.1) by no means fully describes the physics driving melt pool

 $^{^{94}}$ The D86 beam diameter was measured to be approximately 90 μm during the machine maintenance temporally closest to the 1LSB experiments.

formation⁹⁵ but when fitted to experimental data⁹⁶ it is sufficient to inform experiment design decisions. Specifically, the Rosenthal results were combined with the work of Narra [12, Ch. 6] and the balling threshold presented by Francis [31, Ch. 4] to ensure that the test region of process space would include multiple instances of keyholing and balling anomalies. Additionally, the Rosenthal results allowed for the estimation of a maximum expected melt pool width (350 μ m), length (860 μ m), and depth (175 μ m) thus enabling the determination of a conservative⁹⁷ hatch spacing (500 μ m) and melt track length (20 mm) for the 1LSB samples. A custom MATLAB script was developed to perform the Rosenthal analysis described above.

$$T = T_0 + \frac{Q}{2\pi kr} \times \exp\left[-\frac{\nu(r+\xi)}{2\alpha}\right]$$
(6.1)

Where *T* is the temperature as a function of the spatial distance *r* which is defined in (6.3). The term T_0 is the background temperature⁹⁸, *Q* is the absorbed laser beam power, *k* is the thermal conductivity, *v* is the laser beam travel velocity, and α is the thermal diffusivity defined in (6.2). The term ξ is discussed in (6.3).

$$\alpha = \frac{k}{\rho c_n} \tag{6.2}$$

Where k is the thermal conductivity, ρ is the density, and c_p is the heat capacity.

$$r = (\xi^2 + y^2 + z^2)^{1/2}$$
(6.3)
Where ξ is the distance along the beam travel direction with an origin located at the location of the point

Where ξ is the distance along the beam travel direction with an origin located at the location of the point heat source (i.e. the laser beam). The dimensions y and z form a right-hand coordinate system with ξ and can be aligned with the standard coordinate system used throughout this document (Figure 1.6).

⁹⁵ The Rosenthal model assumes steady-state conditions and a point heat source [31]. Additionally, the model neglects temperature-dependent material properties, the latent heat of fusion, and all non-conduction heat transfer mechanisms including fluid flow within the melt pool itself.

⁹⁶ A laser beam absorptivity of 50% and temperature-dependent material properties [228] evaluated at 1427 °C were used to roughly fit the Rosenthal results to the measured melt pool dimensions determined by Narra [12, Ch. 6]. Note that existing work supports the choice of evaluating temperature-dependent material properties near the melting temperature of the material [229]. A preheat temperature of 80 °C was used and the melt pool dimensions were calculated based on the liquidus isotherm of 1336 °C.

⁹⁷ In this context the adjective "conservative" indicates that the single melt tracks are spread apart from each other (spatially) to ensure that clear ex-situ cross-sectional measurements can be made and that they are long enough to ensure steady state melt pool conditions in the cross-sectioned regions. The spatial separation also reduces the chance that a melt track will denude the build surface of powder before the adjacent melt track is exposed.

⁹⁸ In this context, the term "background temperature" refers to the temperature of the material surrounding the melt pool. While this temperature is often directly related to the temperature of the build chamber, it can also be influenced by previous melt tracks [184] and previous layers [224].

Sample Number	Beam Power (W)	Beam Velocity (mm/s)
(EOS Nominal) 1	285	960
2	100	200
3	100	400
4	100	600
5	100	800
6	100	1000
7	150	200
8	150	400
9	150	600
10	150	800
11	150	1000
12	150	1200
13	200	200
14	200	400
15	200	600
16	200	800
17	200	1000
18	200	1200
19	250	200
20	250	400
21	250	600
22	250	800
23	250	1000
24	250	1200
25	250	1400
26	300	400
27	300	600
28	300	800
29	300	1000
30	300	1200
31	300	1400
32	370	400
33	370	800
34	370	1000
35	370	1200
36	370	1400

Table 6.1: Process parameter combinations used for each 1LSB melt track on the EOS M290 L-PBF machine.

To maintain a consistent powder layer thickness for all of the 1LSB melt tracks, the In718 plate was mounted to a modified steel EOS build plate and surface ground. Because the In718 plate had to be removed from the modified steel EOS build plate before being placed inside of the EOS M290, their relative orientations during surface grinding were recorded. As shown in Figure 6.2, layer-wise post-fusion consolidation of the metal powder results in an effective

powder layer thickness that is greater than the thickness of the layers within the additively manufactured part (i.e. the nominal layer thickness). Therefore, in order to better match the nominal 40 µm thick layers used during the in-situ monitoring experiments described in Section 7.2, a 70 µm \pm 20 µm thick powder layer was used for the 1LSB experiment. The uncertainty⁹⁹ in powder layer thickness was quantified by measuring the change in the separation distance between recoater blade and the build plate over the entire area of the build plate; relative gap measurements were performed using a Käfer dial depth gauge. A consolidation factor¹⁰⁰ (\varkappa) of 0.4 (40%) was used in Figure 6.2 based on the density of stainless steel 17-4 PH powder deposited by an L-PBF recoater blade relative to the fused density of the material [191] as reported by Jacob et al. [92]. Note that the true consolidation percentage is dependent upon the particle size distribution of the specific powder system used; however such a measurement is beyond the scope of this thesis. Figure 6.3 shows the powder layer covering the In718 plate during the 1LSB experiments.

 $^{^{99}}$ Systematic error in the powder layer thickness is assumed to be minimal (less than 10 µm) as a result of the procedure used to determine the gap between the recoater blade and the In718 plate: After leveling of the plate with the dial profilometer was complete, the gap was systematically decreased in increments of 1 µm until the recoater blade was unable to travel across the In718 plate. At this point, the plate was dropped 70 µm relative to the recoater blade.

¹⁰⁰ Where the "consolidation factor" is defined as: $\kappa = 1 - \rho_{powder}/\rho_{fused}$, where ρ is the density of the material.





Figure 6.2: A plot showing the convergence of the Figure 6.2: A plot showing the convergence of the Figure 6.2: A plot showing the convergence of the Figure 6.2: A powder layer thickness to a larger, effective of powder layer thickness. Note that the layer-wise a thickness of the deposited material rapidly approaches the nominal powder layer thickness. The data points presented in this plot were artificially generated in MATLAB, i.e. they are extracted from a mathematical relationship¹⁰¹ describing the layer-wise post-fusion powder consolidation process and not from experimental results.

Figure 6.3: An image taken through the EOS viewing window of the powder layer covering the In718 plate after the 1LSB melt tracks were exposed.

6.2.2 Sample Preparation

The 1LSB samples were sectioned (perpendicular the laser beam travel beam direction) using a Wire EDM (Electrical Discharge Machine). The EDM cutting process produces a relatively narrow heat-affected zone [76], reducing the influence of sectioning on measurements of the melt pools. The sectioned samples were then hot-mounted, with the cut face visible, in Buehler[®] Konductomet sample pucks. Each puck contains 10 melt tracks from each of three process parameter combinations. Both halves of the melt tracks were mounted together,

¹⁰¹ The relationship is defined recursively as: $l_{e,i+1} = 0 - H_i$, where $l_{e,i+1}$ is the effective layer thickness at the next layer and H_i is the negative of the height of the build and is given by: $H_i = H_{i-1} + l_{e,i}(1 - \varkappa) - l_n$, where H_{i-1} is the negative of the height of the build at the time of powder spreading, $l_{e,i}$ is the effective thickness of the powder layer spread at the current layer, \varkappa is the consolidation factor, and l_n is the nominal layer thickness of 40 µm. During the first iteration, H_{i-1} is initialized as -40 µm and $l_{e,i}$ is initialized as +40 µm.

therefore 20 cross-sections were available for each process parameter combination. After mounting, each sample was ground and polished according to ASTM E3-11, Table 6 [77]; special care was taken to use an ApexTM Hercules S Pad¹⁰² during the 9 μ m diamond solution polishing. To improve the visibility of the melt pool boundaries, the samples were electro-etched using 10 wt% oxalic acid (C₂H₂O₄) for 15 seconds at 6 V as described by Ramkumar et al. [192]. Note that additional polishing using 0.06 μ m colloidal silica (Buehler[®] MasterMetTM) for 30 seconds to 45 seconds was performed immediately prior to electro-etching in order to the remove the aggressive oxide layer [193] from the samples. Finally, each polished sample was imaged using an Alicona Infinite-Focus optical microscope at an appropriate magnification. The magnifications used for each 1LSB sample are listed in Table 6.2.

6.2.3 1LSB Measurement Techniques

The 1LSB melt pool widths, depths, and cross-sectional areas (Figure 1.7) were manually measured using the Image J software package [78]; Figure 6.4 shows an example micrograph from the 1LSB experiments. Unlike the melt pools in Chapter 2, the full dimensions were measured directly and residual heating was not a concern (see Section 6.2.1). A total of 20 cross-sections were measured for each of the 36 samples (Table 6.1). A selection of 1LSB micrographs and the tabulated melt pool dimension measurements are provided in Appendix E.

¹⁰² Note that Buehler[®] now recommends a TriDentTM Pad for this step when preparing nickel-based superalloys [230].



Figure 6.4: A representative 1LSB micrograph, specifically from Sample #1 (285 W, 960 mm/s). Note the width and depth measurement notations; the cross-sectional area is the region enclosed by the dotted white polygon.

To quantify the uncertainty in the manual measurement of melt pool dimensions, a representative melt pool at each magnification was measured 10 times consecutively. The results of the measurement error are summarized in Table 6.2. The measurement errors range from sample¹⁰³ standard deviations that are 0.35% to 0.75% of the corresponding mean melt pool dimension. The measurement errors are substantially smaller for the In718 melt pools than those reported for AliS10Mg (Table 2.2) owing to their better-defined boundaries post-etching and the ability to measure the full (instead of half) dimensions.

Magnification	Corresponding Sample Numbers	Std. Dev. of Half- Width Measurements (μm, % of mean)	Std. Dev. of Depth Measurements (μm, % of mean)	Std. Dev. of Half- Area Measurements (mm ² , % of mean)
10x	1, 2, 3, 7, 8, 9, 13, 14, 15, 19, 20, 21	0.91, 0.59	1.2, 0.75	6.9×10 ⁻⁵ , 0.45
20x	4, 5, 6, 10, 11, 12, 16, 17, 18, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36	0.55, 0.57	0.64, 0.68	2.3×10 ⁻⁵ , 0.35

Table 6.2: Corresponding measurement errors for the 1LSB experiments.

¹⁰³ In this context "sample" refers to the statistical term "sample of the population" and not the additivelyproduced 1LSB samples.

Based on the 1LSB cross-sections, the 36 power and velocity combinations were each categorized as producing one or more of five different types of melt pools: *desirable, balling,* severe keyholing, keyholing porosity, or under-melting using both qualitative and quantitative measures. Specifically, *balling* melt pools were defined as those exhibiting the characteristic balling morphology [31], [88]. Processing parameters were considered to produce severely keyholed melt pools if the average aspect ratio (defined as the depth over the half-width) was greater than 2.5. Note that in the literature keyhole-mode melting is strictly defined as occurring for any aspect ratio greater than 1.0 as that is the largest aspect ratio achievable under purely conduction-mode melting conditions [31], [83]. Because all of the tested process parameter combinations produced melt pools with an average aspect ratio greater than 1.0 (see Figure 6.8 in the following section), the 2.5 threshold was used as a means of differentiating cross-sectional melt pool morphologies. If any of the melt pool cross-sections for a given power and velocity combination showed examples of keyholing porosity, that process parameter combination was categorized as producing keyholing porosity. Process parameter combinations producing melt pools with average depths less than the powder layer thickness of 70 µm were considered to be *under-melting*. Note that building a bulk part with these parameters would not necessarily result in under-melting porosity as seen in Section 2.3.4 because the melt pool depths may still be greater than the nominal powder layer thickness of 40 µm. For the purposes of Chapters 6 and 7, "under-melting" is used only to indicate that the depth of a melt pool is less than or similar to the depth of the effective powder layer. Finally, any melt pools not otherwise categorized were considered to be *desirable*.

6.3 Results

6.3.1 Melt Pool Geometry

Application of the process mapping approach to the 1LSB samples allowed for the generation of lines of constant melt pool geometries. Specifically, 2D linear interpolation was used to generate a dense matrix of melt pool geometry values (e.g. melt pool width) across beam power and beam travel velocity process space. This dense matrix was then queried such that a set of points in process space was produced at which the relevant melt pool geometry is the same. A smooth curve was then fit to this set of points; in this case a linear function is used; note that this in contrast to the power function used in Section 2.3.1 (refer to that section for a detailed discussion). Table 6.3 shows the R^2 fitting metric between the linear functions and the data for the melt pool geometry process maps. A custom MATLAB script was used to automate the process described above.

Table 6.3: Goodness-o	of-fit metrics for	the In718 melt	pool geometry	process maps.
Table 0.01 000alless (the min for mere	poor Beometry	process maps.

Measurement	R² Value for Linear Fit ¹⁰⁴
Width	0.92
Depth	0.97
Area	0.99
Aspect Ratio	N/A

Figures 6.5 - 6.7 present process maps, respectively, for cross-sectional melt pool width, depth, and area. As expected, the process maps show that higher beam powers and lower beam velocities produce larger melt pools while lower beam powers and higher beam velocities

¹⁰⁴ The reported R^2 values are calculated based the agreement between the model (e.g. linear fit) and the data from each sample (Table 6.1) with a measurement (e.g. melt pool width) within the range presented in the corresponding process map. For example, the reported R^2 values for the melt pool width are based on data from the samples with a measured melt pool width between 100 μ m and 200 μ m (see the legend of Figure 6.5).

result in smaller melt pools. While the melt pool aspect ratios do follow a trend across process space that is similar to the trend observed in Section 2.3.1, the trend is not well-described by simple functional relationships (e.g. linear or exponential fits). Therefore, Figure 6.8 presents the aspect ratio measurements as a 2D linearly-interpolated heat map. The aspect ratio is defined as the depth divided by the half-width, e.g. an aspect ratio of 1.0 indicates a perfectly semicircular melt pool, an aspect ratio less than 1.0 indicates a shallow melt pool, and an aspect ratio greater than 1.0 indicates a deep and narrow melt pool.



·300 μm — 200 μm — 100 μm --50 µm 350 300 (M) 250 200 150 250 100 50 0 200 400 600 1000 1400 1600 800 1200 Beam Velocity (mm/s)

Figure 6.5: Process map of the cross-sectional melt pool **width**, developed from the 1LSB experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the left-right order of the lines of constant geometry matches the left-right order shown in the legend.

Figure 6.6: Process map of the cross-sectional melt pool **depth**, developed from the 1LSB experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the left-right order of the lines of constant geometry matches the left-right order shown in the legend.





Figure 6.7: Process map of the cross-sectional melt pool **area**, developed from the 1LSB experiment data. The error bars represent a 95% confidence interval about the mean. For the reader viewing the figure in grayscale, the left-right order of the lines of constant geometry matches the left-right order shown in the legend.

Figure 6.8: Interpolated heat map of the cross-sectional melt pool **aspect ratio**, derived from the 1LSB experiment data. Note that cross-sectional melt pool aspect ratio may not behave linearly across process space; the use of a heat map to display these data is primarily for visualization purposes.

6.3.2 Distribution of Melt Pool Geometries

As a first step toward understanding the variability of melt pool geometry across process space, the measured size distributions (melt pool cross-sectional width, depth, and area) are shown as cumulative probability plots in Figures 6.9, 6.12, and 6.15. Normalization of the distribution curves was implemented by converting each individual measurement to its percent difference from the mean value for that power-velocity combination. Normal probability plots¹⁰⁵ are shown in Figures 6.10, 6.11, 6.13, 6.14, 6.16, and 6.17 for each geometry measure (cross-sectional melt pool width, depth, and area) for both the EOS nominal process parameter combination and the process parameter combination which deviated the most from a normal distribution. It is evident from both the cumulative probability and normal probability plots that the melt pools which deviate most significantly from the normal distribution form a lower tail. That is, while most of the melt pools follow a normal distribution, in some 1LSB samples several melt pools of a significantly smaller size are present. The implications of this observation are discussed further, and compared to the behavior of the AlSi10Mg alloy, in Section 6.3.6.

To provide context for the use of confidence intervals based on Student's t-distribution [82, p. 419] in the previous section, the melt pool geometry data were quantitatively compared to their equivalent normal distribution. This comparison is shown graphically as the normal probability plots mentioned previously. The 1LSB samples did not provide a sufficient number of measurements to perform a proper Chi-square (χ^2) test¹⁰⁶ [82, Ch. 10]; as a result, Table 6.4 and the legends of Figures 6.9, 6.12, and 6.15 instead present R^2 fit values between each 1LSB data set and its equivalent normal distribution, both of which have been linearized. The majority of the cross-sectional geometry measurements follow normal distributions with the outliers following the trend discussed above.

¹⁰⁵ In a normal probability plot the data are sorted as they would be in a cumulative distribution function and then they are plotted on a non-linear vertical axis representing the normal order statistic medians. If the data are samples which "come from a population with a normal distribution" [227] then they will fall along a straight line [227]. Note that in this context "sample" refers to the statistical term "sample of the population" and not the additively-produced 1LSB samples. Note also that in the implementation [227] used to generate the normal probability plots in this manuscript, the equivalent normal distribution is calculated using only data from the second and third data quartiles.

¹⁰⁶ The standard rule of thumb is that 5 - 8 bins containing a minimum of 5 measurements (i.e. 25 - 40 measurements) are required to perform a valid Chi-square test [82, p. 307].



Figure 6.9: Normalized cumulative probability plots of cross-sectional melt pool **widths** for all 36 1LSB samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.





Figure 6.10: Normal probability plot of the measured **widths** for the EOS nominal process parameter combination (Sample #1). Experimental points far away from the line indicate a deviation from a normal distribution.

Figure 6.11: Normal probability plot of the measured **widths** for the process parameter combination showing the greatest deviation from a normal distribution (Sample #4).



Figure 6.12: Normalized cumulative probability plots of cross-sectional melt pool **depths** for all 36 1LSB samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.





Figure 6.13: Normal probability plot of the measured **depths** for the EOS nominal process parameter combination (Sample #1). Experimental points far away from the line indicate a deviation from a normal distribution.

Figure 6.14: Normal probability plot of the measured **depths** for the process parameter combination showing the greatest deviation from a normal distribution (Sample #28).



Figure 6.15: Normalized cumulative probability plots of cross-sectional melt pool **areas** for all 36 1LSB samples. The five process parameter combinations that produced melt pools with variabilities deviating the most significantly from a normal distribution (see the discussion of R^2 values in this section) are indicated in the legend and with heavier line weights.





Figure 6.16: Normal probability plot of the measured **areas** for the EOS nominal process parameter combination (Sample #1). Experimental points far away from the line indicate a deviation from a normal distribution.

Figure 6.17: Normal probability plot of the measured **areas** for the process parameter combination showing the greatest deviation from a normal distribution (Sample #26).

Table 6.4: Goodness-of-fit between the measured melt pool distributions and their equivalent normal distributions. R^2 values less than 0.80 are highlighted.

Sample Number	R ² (width)	R ² (depth)	R ² (area)
1	0.85	0.92	0.96
2	0.80	0.96	0.94
3	0.95	0.95	0.97
4	0.25	0.93	0.98
5	0.81	0.70	0.65
6	0.98	0.98	0.92
7	0.88	0.84	0.90
8	0.98	0.96	0.97
9	0.92	0.91	0.99
10	0.95	0.96	0.91
11	0.88	0.85	0.92
12	0.72	0.88	0.93
13	0.49	0.91	0.97
14	0.55	0.75	0.96
15	0.97	-0.41	0.12
16	0.73	0.94	0.67
17	0.77	0.57	0.90
18	0.89	0.90	0.90
19	0.92	0.62	0.86
20	0.95	0.86	0.93
21	0.88	0.88	0.96
22	0.95	0.97	0.97
23	0.95	0.94	0.94
24	0.83	0.90	0.87
25	0.94	0.86	0.88
26	0.97	0.03	-13
27	0.93	0.88	0.87
28	0.75	-0.55	-0.65
29	0.94	0.91	0.94
30	0.97	0.78	0.98
31	0.29	0.98	0.96
32	0.86	0.84	0.90
33	0.89	0.64	-0.36
34	0.97	0.95	0.94
35	0.84	0.81	-0.86
36	0.90	0.87	0.97

6.3.3 Variability of Melt Pool Geometry across Process Space

The magnitude of the variability of melt pool geometry across process space was also investigated. Figures 6.18 – 6.20 show the standard deviation (as a percentage of the mean melt pool dimension) for, respectively, the 1LSB cross-sectional melt pool width, depth, and area. The measured variabilities in melt pool width, depth, and area respectively range from approximately 2.9% - 14%, 4.5% - 37%, and 3.8% - 26% of the mean melt pool dimension. In all cases, the melt pool dimension variations are at least one order of magnitude larger than the measurement errors reported in Table 6.2.

The variability, particularly of melt pool depth, is notably higher for the melt pools produced using high beam powers and low beam velocities (the severe keyholing regime of process space), a phenomenon which has been reported previously in the literature for multiple material systems and processes [31], [84]. To a lesser degree, a higher amount of variability, particularly of melt pool area, is also observed for the melt pools produced using the highest beam velocities; this region lies within the balling regime of process space. Section 6.1 provides background on both the keyholing and balling regimes of process space. The aforementioned trends in melt pool variability across process space may be more easily visualized using 2D linearly-interpolated heat maps, such as those provided in Figures 6.21 - 6.23.



Figure 6.18: The variability (standard deviation) in the melt pool **widths** as a percentage of the mean width. Samples are grouped by beam velocity, with each beam power denoted by a different bar hue as shown in the legend. The error bars represent a 95% confidence interval about the sample¹⁰⁷ (percent) standard deviation.



Figure 6.19: The variability (standard deviation) in the melt pool **depths** as a percentage of the mean depth. Samples are grouped by beam velocity, with each beam power denoted by a different bar color as shown in the legend. The error bars represent a 95% confidence interval about the sample¹⁰⁷ (percent) standard deviation.

¹⁰⁷ In this context "sample" refers to the statistical term "sample of the population" and not the additivelyproduced 1LSB samples.



Figure 6.20: The variability (standard deviation) in the melt pool **areas** as a percentage of the mean area. Samples are grouped by beam velocity, with each beam power denoted by a different bar color as shown in the legend. The error bars represent a 95% confidence interval about the sample¹⁰⁷ (percent) standard deviation.




Figure 6.21: Interpolated heat map of the variability (standard deviation) in the melt pool **widths** as a percentage of the mean width.

Figure 6.22: Interpolated heat map of the variability (standard deviation) in the melt pool **depths** as a percentage of the mean depth.



Figure 6.23: Interpolated heat map of the variability (standard deviation) in the melt pool **areas** as a percentage of the mean area.

6.3.4 A Cursory Analysis of the Effects of a Powder Layer on Melt Pool Geometry

As mentioned in Section 6.1, the single bead experiments presented in this chapter were exposed on top of a layer of powder to ensure that the ex-situ data were collected under conditions as close as possible to the in-situ monitoring experiments (Section 7.2) that are the focus of this topic. To demonstrate the importance of considering powder effects under certain conditions (e.g. powder layer thickness, material systems, and processing parameters) the melt pool geometry process maps presented in Section 6.3.1 are compared to the In718 OLSB (single beads without powder) process maps reported by Narra [194, Ch. 6] in Figures 6.24 – 6.26. Note that the OLSB experiments were performed on the same EOS M290 L-PBF machine as the 1LSB experiments presented in this chapter.



Figure 6.24: A comparison between the OLSB (no powder) melt pool **widths** reported by Narra [12, p. 136] (dashed lines) and the 1LSB (with powder) melt pool widths reported in Figure 6.5 (solid lines). For the reader viewing the figure in grayscale, the lines of constant geometry appear in the same left-right order as the top-bottom order shown in the legend.



Figure 6.25: A comparison between the OLSB (no powder) melt pool **depths** reported by Narra [12, p. 137] (dashed lines) and the 1LSB (with powder) melt pool depths reported in Figure 6.6 (solid lines). For the reader viewing the figure in grayscale, the lines of constant geometry appear in the same left-right order as the top-bottom order shown in the legend.



Figure 6.26: A comparison between the OLSB (no powder) melt pool **areas** reported by Narra [12, p. 137] (dashed lines) and the 1LSB (with powder) melt pool areas reported in Figure 6.7 (solid lines). For the reader viewing the figure in grayscale, the lines of constant geometry appear in the same left-right order as the top-bottom order shown in the legend.

While there is reasonably good agreement between the powder versus no-powder lines of constant cross-sectional width, there is significant disagreement between the lines of constant cross-sectional depth and, as a corollary, the lines of constant cross-sectional area. The comparison indicates that in the presence of a powder layer, the melt pools require less beam power to maintain the same size as the melt pools observed in the no powder case. Equivalently, the same beam power and velocity combination produces a larger melt pool in the powder case than in the no-powder case. This result is not unexpected as the dramatically lower thermal conductivity of the powder (relative to the fused material) [100] slows the transfer of heat from the melt pool to its surroundings as mentioned by Montgomery [36].

It is also evident from Table 6.5 that the disagreement is most pronounced for the smaller melt pools with depths less than or similar to the depth of the powder layer, a trend also apparent in simulations reported by Montgomery [36]. Note that more experiments, specifically designed to study this research topic, would be required to robustly quantify the trend. It should be noted that experimental results by Montgomery et al. [195] and Montgomery [36, Ch. 4] indicate that powder layer thicknesses between 20 µm¹⁰⁸ and 70 µm do not measurably influence the melt pool cross-sectional area for the Inconel 625 material system in L-PBF. The reason for this discrepancy in powder layer influence on cross-sectional area is unknown at this time. A direct comparison between measured depths (which saw the largest powder layer dependence) is not possible as melt pool depth values for the no powder case were not reported by Montgomery [36]. Finally, data presented in Section 2.3.4 suggest

¹⁰⁸ Note that building a part with an effective layer thickness of 20 μ m would require a nominal layer thickness of approximately 10 μ m (see the discussion of Figure 6.2) which is not currently standard for any material systems on the EOS M290 L-PBF machine [231] as it would further restrict the usable Powder Systems (see Section 3.2.1).

that the presence of a powder layer may increase melt pool dimensions for the AlSi10Mg material system as well.

Melt Pool Depth (µm)	Melt Pool Depth (percentage of effective layer thickness)	Approximate Effective Power Difference ¹⁰⁹		
50	70%	50%		
100	140%	60%		
200	290%	70%		
300	430%	90%		

Table 6.5: A quantification of the effect of a 70 μ m thick powder layer on melt pool depth for In718.

6.3.5 Melt Pool Morphology Classifications

Figure 6.27 includes examples of all of the melt pool morphology types mentioned in Section 6.2.3. Figure 6.27a shows an example of a *desirable* melt pool cross-sectional morphology. Figure 6.27b shows an example of a *balling* melt pool. Figure 6.27c shows an example of a shallow melt pool that may result in *under-melting*. Figure 6.27d shows an example of a *severely keyholed* melt pool. Figure 6.27e shows an example of a melt pool containing *keyholing porosity*.

¹⁰⁹ In this context, the effective power difference is defined as the ratio of the beam power required to maintain a melt pool of a given depth in the powder case versus the beam power required to maintain a melt pool of the same depth in the no powder case. It is given by: $P_{effective} = 100\% \times (P_{powder}/P_{no \ powder})$.



Figure 6.27: Representative examples of cross-sectional melt pool morphologies.

- (a): A desirable melt pool from Sample #1 (285 W, 960 mm/s.
- (b): A balling melt pool from Sample #35 (370 W, 1200 mm/s).

(c): A under-melting melt pool from Sample #6 (100 W, 1000 mm/s).

(d): A severe keyholing melt pool from Sample #20 (250 W, 400 mm/s).

(e): A melt pool containing keyhole porosity from Sample #7 (150 W, 200 mm/s).

Figure 6.28 shows the distribution of the ex-situ melt pool morphology classifications across process space and will be referred to extensively in Chapter 7. It is worth noting that even the most extreme *keyholing porosity* and *balling* process parameter combinations do not always produce melt pool cross-sections that exhibit those defects. Indeed, for the parameter combinations characterized as *balling*, only between 5% and 85% of the cross-sections exhibited a balling morphology and for the parameter combinations characterized as *poluce* provide the parameter combinations characterized as *balling*, only between 5% and 85% of the cross-sections exhibited a balling morphology and for the parameter combinations characterized as producing *keyholing porosity*, only between 5% and 55% of the cross-sections included keyholing porosity. This is not surprising as both of these flaw formation mechanisms are known in the literature to be periodic in nature (Section 6.1). For this reason it should be expected that these process parameters will produce melt pools of multiple in-situ morphologies imaged by the process

monitoring setup presented in Chapter 7. It is also worth noting that despite all of the tested process parameter combinations producing melt pools with aspect ratios great than 1.0 (indicating keyhole-mode melting [83]), keyholing porosity was only observed in melt pools produced by four of the thirty-six process parameter combinations. As discussed in Section 2.3.4, this suggests the existence of a stable keyholing regime, where even high aspect ratio melt pools are unlikely to generate keyholing porosity.



Figure 6.28: Process space annotated with the ex-situ melt pool morphologies as determined by analysis of the 1LSB cross-sections. The annotations indicate the percentage of melt pool cross-sections which had either keyholing porosity or exhibited the balling morphology. For reference, the EOS nominal parameter combination is indicated and two lines of constant melt pool geometry are overlaid. The line of constant melt pool depth is based on experimental measurements and was calculated as described in the discussion of 6.3.1. The line of constant melt pool length to width ratio is based on the Rosenthal model¹¹⁰.

¹¹⁰ Specifically, the Rosenthal model (6.1) (using the material properties specified in Section 6.2.1) was used to calculate melt pool lengths and widths across beam power and velocity process space. Then lines of constant length over width ratio were calculated as described in Section 6.3.1, except that instead of a linear curve fit, an exponential function of the following form was used: $P = ae^{bv}$, where P is the beam power, v is the beam travel velocity, and a and b are the fitting parameters. A curve was then qualitatively selected such that the length to width ratio dependence of the balling phenomenon is apparent to the reader. Notably, the chosen ratio of 0.24 agrees well with the balling threshold ratio of 0.26 suggested by Yadroitsev et al. [88] and the threshold range of 0.26 – 0.32 determined experimentally (for multiple material systems and AM processes) by Francis [31].

6.3.6 A Comparison of In718 and AlSi10Mg Melt Pool Behaviors

Comparisons between the behavior of the L-PBF-processed In718 discussed in this chapter and the L-PBF-processed AlSi10Mg discussed in Chapter 2 revealed several noteworthy differences and similarities. The noted differences include:

- 1. While curves of constant cross-sectional aspect ratio were found for the AlSi10Mg alloy, equivalent curves could not be identified for the In718 alloy. Specifically, neither linear nor power curves provided fits with reasonably high R² values. Further investigation of other L-PBF-processed alloys would be required to determine whether or not such curves can be identified for the majority or only the minority of material systems. Speculatively, the more complex aspect ratio trends observed for In718 may be caused by the fact that all of the tested parameter combinations lie within the keyholing regime of process space (this is not the case for the AlSi10Mg work), where the dynamics of the vapor pocket may have a strong influence on the cross-sectional aspect ratio.
- The melt pool dimension distribution analyses indicate that occasional deviations from a normal distribution occurred for both alloys, however, in the AlSi10Mg material system the outlier melt pools were larger than average while the outliers were smaller than average in the In718 material system. The precise driver(s) for these outlier melt pools is unknown and warrants further study, however the literature suggests two possibilities:
 (1) Because AlSi10Mg has a relatively low absorptivity (Section 2.1), any transitory changes to the morphology of the keyhole vapor cavity could significantly increase the amount of absorbed beam power (83), [84]. (2) A plausible driver for the smaller In718 outlier melt

pools is the highly dynamic interaction between the vapor plume and the incident laser beam; recent work has suggested that vapor plumes can significantly reduce the amount of beam power absorbed by the substrate [106].

The noted similarities include:

- 1. The magnitude of the observed variabilities in cross-sectional melt pool width and depth (defined as the standard deviation as a percentage of the mean value) were highly similar between the material systems. It should be noted that the measurement error was significantly higher for the AlSi10Mg melt pool dimensions than for the In718 melt pool dimensions (see Sections 2.2.3 and 6.2.3) and this difference may have contributed to the higher measured variability in cross-sectional area for the AlSi10Mg material system compared to the In718 material system.
- 2. Data for both the In718 material system (Section 6.3.4) and the AlSi10Mg material system (Section 2.3.4) suggest that melt pool geometry is substantially altered in the presence of a powder layer. Specifically, the reported data indicate that for a given process parameter combination the melt pool size is larger when exposure occurs on top of a powder layer as opposed to a bare substrate. Critically, this difference can have major consequences for predicted melt pool cross-sectional aspect ratios and lack-of-fusion porosity.
- 3. As discussed in Section 2.3.4, stable keyholing regimes (where keyholing porosity is not present despite high cross-sectional aspect ratios) appear to be present for both material systems. Although MLP experiments, similar to those presented in Chapter

2, for the In718 material system would be required to confirm and better quantify this observation.

6.4 Discussion and Summary

The primary purpose of this chapter is the ex-situ characterization of L-PBF-processed Inconel 718 melt pool morphologies in the presence of a powder layer. These characterizations are used throughout Chapter 7 to enable the interpretation of in-situ melt pool morphologies imaged by a high speed visible-light camera. The data collected for this task also enabled several tangential analyses of interest, the results of which are also summarized here.

In this chapter, correlations between process parameters (beam power and beam travel velocity) and cross-sectional melt pool geometry (width, depth, and area) are presented in the form of process maps for the L-PBF-processed In718 alloy. The correlation between process parameters and aspect ratio is presented as a 2D interpolated heat map because curves of constant aspect ratio were not identifiable. While correlations between process parameters and cross-sectional melt pool geometry have previously been reported for this material system by Narra [12, Ch. 6], those cross-sectional results were based on single measurements of each process parameter combination and the experiments were performed without a powder layer.

Because the process maps presented in this chapter are based on data from multiple melt pool cross-sections at each process parameter combination, the presented process maps have an increased, and quantifiable, level of confidence. Additionally, the relatively large number of cross-sectional measurements allowed for an investigation of the statistical distribution of melt pool dimensions for each process parameter combination. Analysis of the distributions revealed

that cross-sectional melt pool widths, depths, and areas primarily follow a normal distribution with the exception of a handful of outliers (at certain process parameter combinations) which clearly diverge from a normal distribution. Interestingly, the divergent melt pools almost exclusively formed a lower tail, that is, the divergent melt pools were significantly smaller than the majority of the melt pools produced by that process parameter combination. The large sample size also allowed for the quantification of the variability of the melt pool dimensions across process space – critical information for designers as they work at the edges of viable L-PBF processing space. As has been reported for multiple material systems in the literature [31], [84], melt pool depth variability was highest for process parameters within the severe keyholing regime.

The presented process maps for cross-sectional width, depth, and area were compared to their equivalent no-powder process maps reported by Narra [12, Ch. 6]. While there is reasonably good agreement between the powder versus no-powder lines of constant crosssectional width, there is significant disagreement between the lines of constant cross-sectional melt pool depth and area. The comparison indicates that the same beam power and velocity combination produces a larger melt pool in the powder case than in the no-powder case. While this trend has been reported by Montgomery [36], the magnitude of the observed disagreement was substantially higher than reported by Montgomery for L-PBF-processed Inconel 625. The reason for this discrepancy in powder layer influence on cross-sectional area is unknown. It was also observed that the influence of the powder layer increases as the melt pool depth decreases and approaches the thickness of the effective powder layer.

Each process parameter combination was classified based on the size and morphology of its melt pool cross-sections. The melt pool categories are not mutually exclusive and were comprised of: *desirable*, *balling*, *under-melting*, *severe keyholing*, and *keyholing porosity* melt pools. As anticipated, even the most extreme keyholing porosity and balling process parameter combinations did not always produce melt pool cross-sections exhibiting those defects. These ex-situ classifications and defect frequencies are summarized in Figure 6.28 which is referred to extensively in Chapter 7. Finally, the behavior of L-PBF-processed In718 is briefly compared to the behavior of L-PBF-processed AlSi10Mg (Chapter 2) and several noteworthy differences and similarities are identified. One such similarity was the occasional drastic departure of melt pool dimensions from a normal distribution – such unpredictable occurrences further motivate the need for the creation of robust in-situ process monitoring methodologies in order to ensure part quality for critical applications.

- 7 Topic 4: Linking In-Situ and Ex-Situ Melt Pool Morphologies using Machine Learning Techniques in an L-PBF Process
- 7.1 Background and Literature Review

The applications best suited for Additive Manufacturing require a degree of part quality assurance and process reliability that are difficult to achieve with the systems currently on the market [2]. It is commonly recognized that implementation of in-situ process monitoring and closed-loop control is necessary to meet the stringent requirements of these applications [2]. In-situ process monitoring of builds has become a major research topic for the AM community over the last several years. Monitoring efforts for the PBF and DED AM processes have variously focused on detecting macro-scale flaws (e.g. part delamination and residual stress-induced warping) [109], [110], detecting micro-scale flaws (e.g. porosity), measuring temperature fields and histories [109], [110], measuring shielding gas quality [111], and understanding melt pool dynamics [109], [110]. An impressive range of sensor modalities have been explored including those enumerated in the remainder of this paragraph. High speed pyrometers and high speed thermal imaging to measure melt pool temperatures [112]–[114]. Low speed pyrometers and low speed thermal imaging to measure powder bed temperatures [34], [115]–[117]. Embedded thermocouples to measure build substrate temperatures [118]. Low speed visible-light imaging of anomalies on the powder bed [134]-[140], in some cases in conjunction with flash-bulb illumination [141], [142], or structured light (i.e. fringe projection) [143] (see Chapter 4 for a more detailed treatment). High speed X-Ray imaging [85], [125] and interferometric coherence imaging [126] to monitor melt pool sizes and shapes. Strain gages to directly measure part distortion [118], [127]. And perhaps most recently, active [128], [129], passive [130], [131], and spatially resolved acoustic [132], [133] sensing to detect a variety of flaw signatures.

While the powder bed monitoring method presented in Chapter 4 serves an important role, it is fundamentally incapable¹¹¹ of detecting many of the smaller (on the order of 50 µm in size) defects that are of great concern to the AM community. Defects such as porosity formed by the keyholing mechanism [74] and the surface tension-related balling phenomenon [87], [88] are on the same size scale as the melt pool and occur along the melt tracks themselves. Both of these defect types are discussed in detail in Section 6.1 and throughout Chapters 2, 3, and 6. For these reasons, the author considers in-situ monitoring of the melt pool itself to be critical for ensuring part quality.

This view is also shared by many in the AM community, and a substantial body of work now exists relating to the observation of melt pools in L-PBF and DED AM processes using high speed visible-light and thermal imaging. Much of the existing work has focused on monitoring the dimensions of the melt pool. For example, Tan et al. [196] measured melt pool dimensions in a welding process using a coaxially aligned high speed camera and Heigel et al. [123] measured melt pool length in L-PBF using a high speed thermal camera. Impressively, Clijsters et al. [81] developed a real-time system capable of measuring the in-situ melt pool dimensions for an L-PBF process and similar systems are now in use by several L-PBF machine manufacturers [145],

¹¹¹ For a visible-light camera, the physical sensor pixel can be no smaller than the wavelength of visible light (0.4 μ m to 0.7 μ m [205]). As a result, assuming an ideal camera with 1 μ m square pixels, achieving a 10 μ m spatial resolution across the 250 mm × 250 mm EOS build plate would require a 625 MP sensor, at least 25 mm × 25 mm in size. Furthermore, if reliable detection of a flaw requires approximately 10 pixels × 10 pixels of data (a reasonable assumption based on the *patch* sizes used in this chapter) then the sensor size increases to that of the build plate itself.

[146]. Fisher et al. [197] worked to correlate temperature information (collected using a visiblelight camera) with melt pool dimensions in an L-PBF process.

Qualitative and quantitative observation of L-PBF melt pool dynamics has been performed by Gunenthiram et al. [198], Criales et al. [199], and Bertoli et al. [200] although none of these groups sought to directly detect defect formation. Alternate but related sensor modalities have also been used to investigate the melt pool. For example, spectrographic imaging of the vapor plume has been utilized to detect processing defects in the LENS process by Nassar et al. [119]. While Islam et al. [113] attempted to use a combination of a high speed pyrometer and a high speed camera to detect balling via observation of the temperature profile of the melt pool and the surrounding area.

Of most relevance, some work has also applied traditional statistical analysis methods and rudimentary Machine Learning and Computer Vision techniques to the task of defect detection. Luo et al. [201] leveraged a Neural Network to identify the correlation between process parameters and keyhole formation in laser beam welding using data collected from a high speed camera. Grasso et al. (2016) [124] detected off-nominal melting via statistical comparisons between the pixel-wise trans-layer emitted light intensity profiles. Repossini et al. [121] and Grasso et al. (2018) [122] correlated traditional statistical descriptors of spatter (e.g. number and size of spatter particles) and the vapor plume (e.g. emission intensity) with processing parameters (three different energy densities), laser beam scan direction, and catastrophic flaw formation in L-PBF processes. Finally, Khanzadeh et al. [202], [203] presented a method for detecting porosity in the LENS DED process by autonomously clustering different

melt pool morphologies, where the morphologies are defined by a radial function that traces the melt pool boundary as detected by a thermal camera.

The melt pool monitoring approaches reported in the literature have a number of critical limitations: First and foremost, few focus on in-situ flaw detection and many of those that do operate on DED AM processes. In DED, the beam travel velocities are approximately one order of magnitude slower than those used in L-PBF and the melt pool dimensions are approximately one order of magnitude larger than those found in L-PBF (see Section 1.1) – overall DED is a much more conducive environment for in-situ melt pool monitoring than L-PBF. While Khanzadeh et al. [202], [203] utilized unsupervised ML to differentiate between LENS melt pool morphologies, the chosen descriptor of melt pool shape was neither scale invariant nor capable of incorporating information about the spatter or vapor plume. While statistical learning was used by Repossini et al. [121] and Grasso et al. (2018) [122] to characterize spatter and the vapor plume, the work ignored the morphology of the melt pool itself and relied on segmentation rather than more advanced CV *feature* extraction techniques. Additionally, the work of Grasso et al. (2016) [124] requires either comparisons between data collected at different layers or a "clean" signal from a successful build of the same geometry - limiting its applicability to many situations. Finally, none of the work in the literature leverages knowledge of process space to enable the use of supervised ML techniques for melt pool classification and none seek to differentiate between flaw types (i.e. prior work focuses only on detecting "offnominal" melt pools as opposed to specific defect generation mechanisms).

In this chapter, a high speed, off-axis, visible-light camera with a fixed Field of View is used to image Inconel 718 (In718) melt pools in a commercially-available L-PBF process. In this work,

in-situ melt pool morphology is studied and classified using CV and ML techniques in order to identify in-situ flaw formation signatures. It is worth noting that the high speed camera system used in this work detects the thermal emissions from both the melt pool as well as any hot material in the surrounding region (see Section 7.2.2). Therefore the in-situ data are not well predicted by existing simulations - motivating the usage of CV and ML techniques to describe these data. Because the melt pool size often does not correlate to ex-situ flaws, a scale-agnostic description of melt pool morphology is created by applying the Bag-of-Words (or Keypoints) [44] ML technique to features extracted using SIFT [204]. The presented description of melt pool morphology explicitly incorporates information regarding the shape of the melt pool itself, the vapor plume, and spatter in order to improve differentiation between in-situ morphologies. Acquiring ground-truth information about in-situ melt pool morphology in a commerciallyavailable L-PBF system is extremely non-trivial. Therefore, the ex-situ melt pool morphology data collected in Chapter 6, combined with fundamental knowledge of process space [29], are used to bridge the gap between unsupervised and supervised Machine Learning. This novel, human-in-the-loop, ML approach enables not only the identification of similar and dissimilar melt pool morphologies, but also the preliminary classification of melt pools based on observed in-situ flaw signatures. Specifically, the presented ML methodology is capable of classifying melt pool morphologies into four categories: desirable, balling, under-melting, and keyholing porosity which are defined in Chapter 6 as well as an additional category referred to as spatter which is defined in Section 7.3.6.

Furthermore, while the work presented in Chapters 2 and 6 demonstrates that process parameters can be chosen to reduce the likelihood of melt pool-scale defects, the approach is

not easily extensible to non-bulk geometries. In other words, even if process parameters are chosen to reside within a desirable processing window based on bulk data (e.g. Figures 2.35 and 6.28), defects may still occur during the printing of certain geometries such as narrowing stripes (Figure 7.1a), thin wall structures [183], unsupported overhangs [182] (Figure 7.1b), and contours (Figure 7.1c). As a result, the only way to assure machine users that melt pool-scale flaws are not occurring is through the implementation of in-situ process monitoring or prohibitively-expensive ex-situ testing [52]. To address this issue, the high speed camera is used to collect data during the printing of test artifacts containing the four non-bulk geometries listed above. The imaged melt pools are then analyzed using the trained ML methodology and flaw formation triggered by local build geometry is observed for some test artifacts.



Figure 7.1: Simplified representations of three different non-bulk geometries. (a) A stripe decreasing in width as it intersects the edge of a part at an angle. As the width decreases, residual heating from adjacent melt tracks [184] is expected to influence the morphology of the melt pool. (b) An unsupported overhang. As the melt pool passes over the bed of un-fused powder, the differing thermal conditions [36], [100] are expected to influence the morphology of the melt pool. (c) A contour pass conforming to the exterior geometry of a part. The differing thermal conditions near the edge of the part [100] are expected to affect the morphology of the melt pool.

Finally, the high data burden associated with the real-time monitoring of the melt pool

throughout an entire build is not ignored¹¹², and possible, albeit not immediate, real-time solutions are briefly discussed. The work presented in the chapter was supported by CMU's Manufacturing Futures Initiative (internal grant number 062900.005.105.100020.01) and the purchase of the high speed camera and associated optics was supported by a Carnegie Institute of Technology Dean's Equipment Grant, FY 2016.

 $^{^{112}}$ At 6,400 fps and a resolution of 1024 pixels × 1024 pixels a high speed camera would produce approximately 13 Gb of data per second. As a result, over the course of a relatively short, 24 hour L-PBF build, on the order of 1000 Tb of data would be generated. Such a data burden is currently unmanageable both in terms of data storage capacity as well as data transmission times.

7.2 Experimental Design and Methods

7.2.1 Programming Environment

Unless otherwise noted, all software was developed within the MATLAB R2016a or R2017a environments. The above MATLAB versions also included the following add-on packages: the Image Processing Toolbox and the Statistics Toolbox.

7.2.2 High Speed Camera

The in-situ melt pool images analyzed in this thesis were captured by a Photron FASTCAM Mini AX200 high speed camera mounted to the EOS M290 at CMU's NextManufacturing Center as shown in Figures 1.4 and 7.2. The author would like to note that this setup was developed and characterized by Brian Fisher [34] at CMU; as such, all of the information presented in this subsection is for reference only. The light captured by the camera sensor is approximately within the visible range and originates as thermal emissions¹¹³ from the melt pool and surrounding material. The high melt pool temperatures (Section 6.1) allow for sufficient emission in the visible spectrum [205] for the use of a visible-light camera to be practical. Additionally, traditional thermal cameras (sensitive to frequencies of light in the infrared [205]) are typically more expensive and limited to lower frame rates than visible-light cameras.

The high speed camera is equipped with a 1024 pixels \times 1024 pixels (1 MP) sensor with a depth (dynamic range) of 12 bits. It is capable of maximum recording speeds of between 6,400

¹¹³ It is worth reiterating that this implies that the true melt pool boundaries are unknown when viewing the melt pool images. Further characterization of the system is explored by Fisher [34] which enables the estimation of melt pool size under certain conditions. Nonetheless, the inability to precisely identify the true melt pool boundary is not a hindrance to the presented study of melt pool morphology.

fps and 900,000 fps, depending on the percentage of active pixels, and an exposure time as low as 260 ns. Data are recorded on an onboard 16 GB circular RAM buffer. For the experiments presented in this chapter, camera recording was triggered¹¹⁴ either manually by the author or automatically by the Photron PFV software package [206] based on the sudden increase in brightness corresponding to a melt pool entering the Field of View (FoV). The optical train that is used in this work provides a maximum field of view of 6.35 mm × 6.35 mm, with each pixel covering a 6.2 μ m × 6.2 μ m area. The resolving power of the system is approximately 50 linepairs per millimeter as determined by visual inspection of a negative 1951 USAF resolution test target supplied by Thorlabs.



Figure 7.2: The high speed camera setup mounted to CMU's EOS M290; Figure 1.4 shows this setup schematically. This image of the high speed camera was taken by Brian Fisher of CMU.

¹¹⁴ Strictly, the trigger point actually stops the recording, preventing the information stored in the circular buffer from being overwritten.

7.2.3 Camera Calibration Procedure

The location of the available chamber viewport (Figure 7.2) necessitates that the high speed camera be mounted such that its axis is not parallel to the normal vector of the build plate (*z*-axis). The resulting distortion is corrected using a fully-constrained Homography matrix [154] to apply an affine warp to the raw image such that a rectangular object in the initial image will appear rectangular (in the correct proportions) in the final image. In order to determine the appropriate Homography matrix, a rectangular fiducial exposed using the nominal EOS parameters (Table 6.1) was imaged using both the high speed camera (Figure 7.3) and an Alicona Infinite-Focus optical microscope (Figure 7.4). Specifically, the four interior corners of the fiducial were used to map the high speed camera image to the "true" dimensions as measured from the microscope image. As expected, the camera distortion was minimal, given the relatively small FoV and shallow camera angle (relative to the *z*-axis). A similar procedure is also described in Section 4.2.3.



Figure 7.3: The rectangular fiducial as imaged by the high speed camera.



Figure 7.4: The rectangular fiducial as imaged by the Alicona Infinite-Focus microscope.

The frame rate of the high speed camera was chosen a priori to be 6,400 fps – the maximum frame rate for which the entire FoV (1024 pixels \times 1024 pixels) can be captured. At this frame capture rate the 16 GB RAM buffer on the high speed camera can record data for approximately 1.7 seconds. An exposure time of 5.00 µs was chosen based on visual observation of melt pools produced using four different power and velocity combinations spanning the process space of interest. Specifically, melt pools produced using PV #1 (nominal EOS parameters), PV #19 (high energy density regime), PV #6 (low energy density regime), and PV #36 (balling regime) were imaged at multiple exposure times (see Table 6.1). The final exposure time was chosen such that the low energy density melt pool was visible while the vapor plume of high energy density melt pool did not completely obscure the melt pool itself. During a 5.00 μ s exposure time a melt pool traveling at the maximum tested velocity of 1400 mm/s will travel 7 µm, or approximately one pixel; therefore motion blur can be considered negligible under these conditions. At a frame rate of 6,400 fps the melt pool will travel between 30 µm and 220 µm (over the range of beam velocities tested) between camera frames. These inter-frame travel distances are comparable to the length of the melt pool (Section 6.2.1). Brian Fisher of CMU lent their significant expertise in this content area to help the author determine the appropriate camera settings.

7.2.4 Coaxial Melt Pool Transformation

The high speed camera configuration described in Section 7.2.1 provides a stationary FoV; as a result, the melt pool data are collected in an Eulerian frame of reference. This increases the difficulty of comparing data between melt pool images, and for this reason a Lagrangian frame

of reference is preferred. While this can be accomplished via a coaxially-aligned optical configuration as demonstrated by Clijsters et al. [81] and Lane et al. [120], such an implementation is extremely non-trivial for an AM machine user that is not also an AM machine manufacturer. Therefore, in this chapter, a custom software solution is pursued that is capable of transforming the collected melt pool data into data that appears to have been collected via coaxially-aligned imaging. Figure 7.5 shows a representative image collected by the fixed FoV high speed camera. Figure 7.6 shows the same image after the melt pool has been transformed into the coaxial (Lagrangian) reference frame.





Figure 7.5: A representative melt pool image from the high speed camera with a fixed (Eulerian) reference frame. The approximate travel path of the melt pool is indicated by the dashed lines; the spacing between the dashed lines is not necessarily equal to the hatch spacing. The false-color intensity map was chosen to increase the visibility of the data given its spread across a relatively wide dynamic range.

Figure 7.6: The melt pool image from Figure 7.5 after transformation. The reference frame is coaxial, i.e. Lagrangian. The nose, tail, and centerline (dotted line) of the melt pool are indicated. High speed camera data are first warped using the Homography matrix described in the previous subsection (7.2.3). Then each frame is binarized using a manually-determined intensity threshold intended to remove (only) sensor noise. A connected-components [207] analysis is performed to differentiate the melt pool itself from other bright objects such as spatter [208]. All but the largest connected component are removed to prevent spatter from influencing the transformation calculations. The first transformation parameter calculated is the in-plane rotation angle; this task is accomplished by locating the centroid of the melt pool in Figure 7.5 and comparing its relative motion with respect to the centroid of the melt pool in a subsequent camera frame. The stability of this calculation is further improved by additional thresholding such that only the brightest region of the melt pool is considered.

Once the in-plane rotation has been determined, it is applied to the binarized (noise threshold) melt pool image (with spatter removed). At this stage the "nose" (Figure 7.6) of the melt pool is identified and used to place the melt pool within a consistent, coaxially viewed bounding box. The nose coordinates are defined as the centroid of the melt pool along the axis perpendicular to the travel direction (i.e. the line of symmetry of the melt pool) and the leading edge of the melt pool along the axis parallel to the travel direction (i.e. the ξ -axis). The choice of using the nose to define the melt pool position was made due to the relative temporal stability of the nose, compared to the melt pool centroid, which is affected by the highly variable melt pool tail (Figures 7.6 and 7.7).

The in-plane rotation and bounding box transformation parameters are now applied to the original camera data to isolate the melt pool. Figure 7.7 shows subsequent frames taken by the high speed camera. Each frame underwent the coordinate transform described above such that

the melt pool in each frame appears in the same spatial location as the melt pool in the prior frame. At this point additional operations can be applied, e.g. the spatter can be removed (as described above), the spatter can be isolated (by removing the largest connected component), thresholds designed to extract the vapor plume can be used, etc. While sudden changes in melt pool travel direction are handled naturally by the described methodology, special care must be taken if the melt pool passes out of the FoV of the camera entirely or "disappears" as is common if skywriting [209] is implemented at the edges of raster melt tracks. Design of appropriate heuristics for these situations is beyond the scope of this thesis and such cases were handled manually.



Figure 7.7: Eight subsequent high speed camera frames, transformed such that the melt pool appears to be spatially stationary and varying temporally. The relative size scale is preserved between all of the images.

7.2.5 Experiment Sequence Overview

Data used in this chapter were collected over the course of five different sets of experiments (not counting the previously described Homography fiducial). As discussed in Section 7.1, information regarding ex-situ melt pool morphology is required to enable the linkage of in-situ melt pool morphology to processing defects. The necessary ex-situ data were collected and analyzed in Chapter 6 which is collectively referred to as Experiment 1 (E1) throughout the rest of this chapter. During Experiment 2 (E2), the high speed camera was used

to image melt pools produced by 36 different process parameter combinations spanning the process space of interest. These data were then used to train the ML algorithm presented in Section 7.3. Finally, Experiments 3a - 3c (E3a - E3c) investigated melt pool morphology during exposure of non-bulk geometries. Specifically, the high speed camera was used to image melt pools as the laser beam approached the edge of a stripe (E3a), as the laser beam passed over an unsupported overhang (E3b), and as the laser beam exposed contours (i.e. traveled along the edge of a part) (E3c). Experiments E2, and E3a - E3c are described in the following two subsections. All of the experiments were performed on an EOS M290 L-PBF machine at CMU's NextManufacturing Center and utilized low carbon steel build plates of size 1.5 in \times 6 in (visible in Figures 7.8 and 7.9). The use of these sub-size build plates allowed for a decrease in turnaround time between the E2 and E3 experiments.

7.2.6 Collection of In-Situ Training Data (Experiment 2)

As described in Section 7.2.2, the high speed camera is only able to image a single, fixed FoV; therefore in-situ data collection during E1 could not be performed. To overcome this challenge, the same 36 process parameter combinations used in E1 and enumerated in Table 6.1 were exposed within the FoV of the high speed camera during E2. Specifically, for each parameter combination a set of 10 mm long $1LSB^{115}$ melt tracks were exposed within a rectangular region¹¹⁶ of size 10 mm × 20 mm. Each melt track was separated by a hatch spacing

¹¹⁵ One Layer Single Bead (1LSB) experiments, i.e. one layer of powder, single bead exposures.

¹¹⁶ The width (20 mm) of the rectangular exposure area was designed to be substantially larger than the width of the FoV (6.2 mm) because it was not known a priori how challenging it would be to align the high speed camera such that its FoV would cover the region of interest. Future experiments of this type can make use of a significantly reduced alignment safety margin.

of 500 μ m and adjacent melt tracks were not exposed subsequently i.e. at least 140 ms elapsed between the exposures of adjacent melt tracks. As in E1, the chamber preheat was 80 °C, the nominal¹¹⁷ beam diameter was 100 μ m, and the nominal powder layer thickness was 40 μ m (resulting in an approximately 70 μ m effective layer thickness, see Section 6.2.1).

After the exposure of the 1LSB melt tracks for a given parameter combination, the rectangular region was re-exposed (without the addition of another powder layer) using the EOS nominal parameters (PV #1) and the nominal hatch spacing of 110 µm. Then, an additional 11 layers of material were built using the EOS nominal parameters such that each set of 1LSB melt tracks were vertically separated by a 440 µm tall block of nominally-processed material. Vertical separation was implemented to ensure that the observed melt pool morphologies were not influenced by an interaction with porosity (or other defects) left behind by previously tested process parameter combinations. The initial set of 1LSB melt tracks were separated from the sub-size steel build plate¹¹⁸ by approximately 1 mm of nominally-processed material. Figure 7.8 shows a Computer Aided Design (CAD) rendering of the E2 experiment and Figure 7.9 shows the as-built E2 block of material.

 $^{^{117}}$ The D86 beam diameter was measured to be approximately 90 μm during the machine maintenance temporally closest to the E2 experiments.

¹¹⁸ A sub-size low-carbon steel plate (McMaster-Carr P/N: 1388K311) mounted to a modified EOS build plate was used to decrease the turnaround time between the sequence of experiments described in Section 7.2.5.





Figure 7.8: A CAD rendering of the E2 rectangular exposure region built up over many sets of 1LSB melt tracks and separation layers. Note the FoV of the high speed camera.

Figure 7.8: A CAD rendering of the E2 rectangular Figure 7.9: The as-built material deposited over the exposure region built up over many sets of 1LSB melt course of E2.

The 10 mm length of each 1LSB melt track was chosen such that the melt pool was always a minimum of 2 mm away from the end of the melt track while it was in view of the camera. More discussion of steady state melt pool lengths is available in Section 6.2.1. The hatch spacing was determined as described in Section 6.2.1 and the vertical buffer distance between sets of 1LSB melt tracks was informed by the maximum expected melt pool depth¹¹⁹ of 175 µm reported in Section 6.2.1. To support robust training of the ML methodology, a substantial amount of training data are required – a minimum of 500 camera frames for each of the 36 power and velocity combinations was deemed sufficient for this task based on prior ML work by the author (Sections 4.3.2 and 4.4.2). The approximate number of frames expected to be captured by the camera per set of 1LSB melt tracks can be calculated using (7.1) – (7.4) as discussed below.

¹¹⁹ The buffer depth was chosen based on the maximum predicted melt pool depth of 175 μ m and the depths reported by Narra [12, Ch. 6]. The maximum measured depth of any melt pool was 600 μ m. This only barely exceeds the sum of the vertical separation depth (440 μ m) and the penetration depth of the melt pools produced with the nominal parameters (150 μ m).

$$N = \left(\frac{L_{fov}}{L_{tot}}\right) t_{data} f \tag{7.1}$$

Where *N* is the number of recorded camera frames, L_{tot} is the total length of the exposure rectangle (10 mm), L_{fov} is the length of the melt track within the FoV (6.2 mm), t_{data} is the data collection time given by (7.2), and *f* is the camera frame rate (6,400 fps).

 $t_{data} = \min[t_{inframe}, t_{maxrecord} - t_{tot}]$ Where $t_{inframe}$ is the amount of time the melt pool spends within the FoV given by (7.3) and $t_{maxrecord}$ is the maximum recording time for the camera (1.7 seconds). While not the case for any of the tested process parameter combinations, special care must be taken if t_{tot} exceeds $t_{maxrecord}$. In the case of $(t_{maxrecord} - t_{tot}) < t_{inframe}$ a conservative t_{data} estimate is produced. (7.2)

$$t_{tot} = \left(\frac{L_{tot}}{v}\right) \left(\frac{W_{tot}}{h}\right) \tag{7.3}$$

Where t_{tot} is the total time it takes expose the entire rectangular region, v is the laser beam velocity, W_{tot} is the total width of the exposure rectangle (20 mm), and h is the hatch spacing (0.5 mm).

$$t_{inframe} = \left(\frac{L_{tot}}{v}\right) \left(\frac{W_{fov}}{h}\right) \tag{7.4}$$

All of the variables have been previously defined.

Using the equations above, it was determined that for several of the process parameter combinations more than one set of 1LSB melt tracks would be required to generate the required 500 frames of training data. The observant reader may notice that the E2 1LSB melt track length (10 mm) is shorter than that used for E1 (20 mm). This reduction in track length was motivated by the need to reduce the total experiment time which was substantially dependent upon the number of sets of 1LSB melt tracks required for each parameter combination¹²⁰.

The number of frames of training data collected for a single process parameter combination ranged from 504 to 1,394. A total of 24,484 frames of usable data were collected across 29 of the 36 process parameter combinations. Data from four of the parameter

¹²⁰ As can be seen in (7.2) and (7.3), if the total length of the melt track is too long (L_{tot}) the maximum camera recording time ($t_{maxrecord}$) will dominate the time for which data can be collected (t_{data}) thereby reducing the number of frames collected for each set of 1LSB melt tracks (N).

combinations in the low energy density regime (PV #5, #6, #12, and #18) were not collected as the apparent melt pool size (brightness) was too small to automatically trigger the high speed camera's recording feature. Data from a further three parameter combinations (PV #27, 28, and #29) were inadvertently recorded using an incorrect exposure time and were therefore not used for training. A selection of high speed camera images from across process space is available in Appendix F.

7.2.7 Non-Bulk Geometries (Experiments 3a – 3c)

Experiment 3a was designed to investigate melt pool behavior near the edges of a stripe. As shown in the CAD rendering of the test artifact in Figure 7.10, the FoV is centered on the boundary of a rectangular region of size 10 mm \times 20 mm. In other words, this configuration allowed imaging of melt pools as they approached (and departed from) a stripe boundary with the specific hope of observing any morphology changes triggered by local preheating from adjacent melt tracks [184]. In order to determine the effect of stripe width on near-edge morphology, data were collected for several stripe widths ranging from 10 mm to 0.5 mm. Before data collection (and stripe width reduction) began, the test artifact was vertically separated from the sub-size build plate by a block of nominally-processed material (PV #1 with a hatch spacing of 110 μ m) approximately 400 μ m in height.

Finally, a thin wall structure of size 0.5 mm \times 20 mm was built to a height of 5 mm while in-situ imaging data were collected at wall-height intervals of 1 mm. Note that this thin wall was exposed using a sequence of adjacent melt tracks running the length of the short dimension (as shown for the stripes in Figure 7.10) while thin walls are more typically exposed using adjacent

melt tracks running the length of the long dimension. The E3a test artifact was exposed with the EOS nominal parameters (PV #1) and a hatch spacing of 110 μ m. As in E1 and E2, the chamber preheat was 80 °C, the nominal beam diameter was 100 μ m, and the nominal powder layer thickness was 40 μ m. Figure 7.11 shows the as-built test artifact after completion of E3a. Over the course of E3a, at least 30 frames of data were collected for each stripe width.



Figure 7.10: A CAD rendering of the completed E3a testFigure 7.11: The as-built test artifact deposited over
the course of E3a.

Experiment 3b was designed to investigate melt pool behavior during the fusion of an unsupported overhang. Specially, the goal was to observe any melt pool morphological changes triggered by the differing thermal conditions present due to the low thermal conductivity of the unfused powder [100], [210] below the overhang. As shown in Figure 7.12, the FoV is centered on an overhang spanning a gap of 5 mm. To decrease the turnaround time between the experiments (Section 7.2.5) a channel was milled into several of the sub-size plate – thereby reducing the build time and allowing additional E3b test artifacts to be fabricated in the event of a build or data capture failure. Pads 160 µm in height¹²¹ were built using nominal parameters

¹²¹ Note that this height exceeds the predicted PV #1 melt pool depth of 70 μ m and is slightly greater than the average measured depth of 150 μ m.

(PV #1) and a hatch spacing of 110 μ m on either side of the pre-cut channel in order to ensure appropriate bonding between the overhanging layers of interest and the sub-size plate.

Data were collected for a total of 5 layers spanning the channel. The first observed layer was exposed directly on top of a bed of unfused powder while the fifth layer was exposed on top of (nominally) 160 μ m of fused material. The E3b test artifact was exposed with the EOS nominal parameters (PV #1) and a hatch spacing of 110 μ m. As in E1 and E2, the chamber preheat was 80 °C, the nominal beam diameter was 100 μ m, and the nominal powder layer thickness was 40 μ m. Figure 7.13 shows the as-built test artifact after completion of E3b. Over the course of E3b, approximately 1,500 frames of data were collected for each of the five overhanging layers.



Figure 7.12: A CAD rendering of the completed E3b testFigure 7.13: The as-built test artifact deposited over
the course of E3b.

Experiment 3c was designed to investigate melt pool behavior during the exposure of contours. Specially, the goal was to observe any melt pool morphological changes triggered by the differing thermal conditions present due to the low thermal conductivity of the unfused powder [100], [210] surrounding the part. As shown in Figure 7.14, the FoV is centered on the

edge of a rectangular test artifact of size 20 mm \times 10 mm such that the melt pool could be imaged as it traveled along the edge of the test artifact. The E3c contours were exposed with two different sets of process parameters enumerated in Table 7.1. The first set of contour parameters are simply the EOS nominal¹²² bulk parameters (PV #1) while the second set of parameters were designed by the author and Dr. Sneha Prabha Narra of CMU to mitigate nearsurface keyholing porosity¹²³. The offset distances between the first set of contours and the edge of the rectangular test artifact are half of the predicted (Section 6.2.1) melt pool width (140 µm for PV #1).

Table 7.1: Process parameter	combinations used	for each set of	f contours imaged	during E3c.
------------------------------	-------------------	-----------------	-------------------	-------------

	Contour Parameter Set #1										
Outer Contour			Middle Contour		Inner Contour						
	Beam Power (W)	Beam Velocity (mm/s)	Distance from Part Edge (µm)	Beam Power (W)	Beam Velocity (mm/s)	Distance from Part Edge (µm)	Beam Power (W)	Beam Velocity (mm/s)	Distance from Part Edge (µm)		
	285	960	0	285	960	75	285	960	150		
Contour Parameter Set #2											
	110	300	0	125	300	70	125	300	140		

The initial set of E3c contour melt tracks were vertically separated from the sub-size steel build plate by a block of nominally-processed material (PV #1 with a hatch spacing of 110 μ m) approximately 1 mm in height. Both sets of contour parameter combinations were exposed and imaged twice in order to collect additional data. As in E2, each set of exposures were vertically separated with 440 μ m of nominally-processed material. As in E1 and E2, the chamber preheat was 80 °C, the nominal beam diameter was 100 μ m, and the nominal powder layer thickness

¹²² EOS nominal contour parameter combination for In718 is 138 W, 300 mm/s.

¹²³ This contour development work was performed for General Electric as part of a project through CMU's NextManufacturing Center.

was 40 µm. Figure 7.15 shows the as-built test artifact after completion of E3c. Over the course of E3c, at least 65 frames of data were collected for each contour.





Figure 7.14: A CAD rendering of the completed E3c test Figure 7.15: The as-built test artifact deposited over artifact. Where Δ is the distance from the contour melt track to the edge of the part. Note the FoV of the high speed camera.

the course of E3c.

7.3 Bag of Words (BoW) Methodology and Theory

7.3.1 Overview

The methodology presented in this section is an application of a widely-used ML technique, known as Bag of Words (or Keypoints) (BoW) [44], often applied to CV problems. In this implementation, the training data consist of frames of data captured by the high speed camera during E2 and transformed such that the melt pool appears to be in a Lagrangian (coaxial) reference frame (Section 7.2.3). The only human-applied labels associated with the training data are the process parameter combinations used to produce each observed melt pool. While the BoW technique can be applied to multiple *feature* types, the author chose to use SIFT features for their ability to be agnostic to scale information. Such agnosticism was deemed important as melt pool size does not necessarily correlate with the ex-situ flaws identified in

Chapter 6. This section is intended to provide an overview of this methodology along with relevant ML and CV theory. Figure 7.16 is a flowchart of this ML methodology and is referred to extensively throughout this section.



Figure 7.16: Flowchart of the implementation of the BoW ML technique discussed in this section. The representative micrographs shown in step (i) are also shown in Figure 6.27. The human "stick figure" shown in between steps (h) and (i) was inspired by the XKCD web comic series created by Randall Munroe [211].

7.3.2 Selection of the Training Data

Training data were collected during E2 as described in Section 7.2.6 – the high speed camera was used to image melt pools produced using a total of 29 different process parameter combinations spanning the EOS M290 process space. Each of the 24,484 frames of usable data were warped and coaxially transformed as described in Sections 7.2.3 and 7.2.4. To improve the
robustness of the coaxial transformation during construction of the training database, the appropriate in-plane rotation angle was determined a priori using a subset of training images and then applied to all of the subsequent training images. Unlike in Sections 4.3.2 and 4.4.2, the morphology of each training image is not labeled with a "ground truth" classification by the author. Such an action is impossible as the correlations between the in-situ appearance of a melt pool and ex-situ outcomes are not known. Indeed, this challenge necessitated the development of the ex-situ database in Chapter 6 and the identification of linkages between the in-situ and ex-situ melt pool image only indicate the process parameter combination used to generate the corresponding melt pool; the usage of these labels is discussed in Section 7.3.6. The final training database is composed of a total of 24,385 coaxially-transformed melt pool images labeled with their associated process parameters.

7.3.3 Scale Invariant Feature Transform (SIFT)

This implementation of the BoW ML technique extracts *features* using the SIFT algorithm. SIFT *features*, as their name (Scale Invariant Feature Transform) might suggest, are considered robust even when the objects of interest in an image (or across images) may be of varying sizes. First developed by Lowe [204], they are commonly used when it is an object's overall shape that is of interest. The potential for scale agnosticism is considered by the author to be critical, as melt pool size often does not correlate to ex-situ flaws. For example, two different process parameter combinations may produce melt pools with similar widths, however their morphologies could be radically different with one melt pool considered *desirable* and the other prone to generating *keyholing porosity*. Interestingly, this is not SIFT's first foray into the AM community as DeCost et al. [212] have used SIFT to classify metal powders for AM applications and a conceptually-similar *feature* known as DAISY has been used by Jacobsmühlen et al. [140] to detect *super-elevation*.

SIFT *features* characterize the gradient field surrounding each pixel in an image. In this work, gradient orientations within a 2 pixels × 2 pixels window are considered. That is, a coaxially-transformed melt pool image (Section 7.3.2) is broken into non-overlapping windows (Figure 7.16b), each containing 4 pixels. The gradient orientations are grouped into nine, unsigned bins, e.g. one of the bins encompasses all gradients with the following orientations: 0° $- 20^{\circ}$ and $180^{\circ} - 200^{\circ}$. Note that unlike the *filter* convolution operation (with stride of 1) described in Section 4.3.3, the output of this process is an image with a spatial resolution that has been reduced by a factor of two in each direction. While it is common practice to only calculate SIFT *features* at interest points¹²⁴ (a.k.a. keypoints) in an image [213], in this implementation, a dense field of SIFT *features* is calculated at every strong edge. Where a "strong edge" is defined as any pixel at which the magnitude of the gradient is at least 10% of the magnitude of the strongest gradient in the image. Figure 7.18 shows a visualization of the SIFT *descriptors* applied to a coaxially-transformed melt pool image (Figure 7.17).

¹²⁴ Corners are commonly used as keypoints because they can be detected robustly in most "every-day" images [213]. Common applications of *feature* extraction at corners include panorama stitching and object matching within a scene.





Figure 7.17: An example false-color and coaxiallytransformed image of a melt pool produced using PV #36.

Figure 7.18: The SIFT *features* extracted from the melt pool shown in Figure 7.17. This visualization of the SIFT *features* was generated using a MATLAB script written by Dollar [214].

7.3.4 Histogram of Oriented Gradients (HOG)

SIFT *features* are traditionally considered highly specific, that is, they have a dimensionality equal to the number of orientation bins (i.e. nine) with values in each dimension spanning a subset¹²⁵ of \mathbb{R} . For this reason, SIFT algorithms are often used for template matching applications [213] and less often for classification. A variety of techniques are available to artificially reduce the specificity¹²⁶ of the SIFT *features*. Additionally, although not pursued in this chapter, the dimensionality of a *feature* can also be reduced using methods such as PCA (Principle Component Analysis) [215]. For this work, specificity reduction was accomplished using an unsupervised clustering approach similar to that described in Section 4.3.4.

¹²⁵ Only numbers of the set $\mathbb{Z}[0, W]/W$ are possible, where W is the window size (i.e. four pixels).

¹²⁶ A simple method for specificity reduction is rounding a given value to a set number of digits after the decimal point. For example, reducing the \mathbb{R} space to the subspace of $\mathbb{Z}/10$ would eliminate the distinction between the values 42.14 and 42.13.

In order to utilize a pre-existing high-dimensional clustering method, the SIFT *features* were first converted into a standard vectorized format known as Histogram of Oriented Gradients (HOG) [216] (Figure 7.16c). Specifically, the number of gradients within the SIFT window falling into each orientation bin is counted and stored in the corresponding element in the HOG vector (i.e. for nine orientation bins the HOG vector will be nine elements long). The values in each element of the HOG vector are then normalized such that they range in magnitude from \mathbb{R} [0, 1]. SIFT *features* (and their HOG equivalents) were collected from all of the melt pool images in the training database. Note that these *features* were extracted from each training image under three different contrast adjustments which are further discussed in the following subsection. Note that no subsampling of the training data is performed, i.e. all of the extracted SIFT *features* are included in the training process.

Once collected, HOG vectors with similar values in each element (i.e. vectors that describe a similar gradient field) are grouped together using a standard k-means unsupervised clustering algorithm [163], represented by Figure 7.16d. For this work, cluster initialization was performed using random seeding, with preference given to a uniform spacing between clusters. During development of this ML methodology, the requested number of clusters was varied between 25 and 200; satisfactory performance was achieved with 50 clusters. Each cluster is represented by a mean HOG vector. The 50 mean HOG vectors are commonly referred to as *visual words*, and are stored in a *dictionary*, represented by Figure 7.16e.

After the *dictionary* is constructed, each HOG vector in each training image can be matched to the closest (pair-wise distance [164]) *visual word* in the *dictionary* (Figure 7.16f). Note that this operation has reduced the set of possible values of the SIFT *features* to the set

251

 $\mathbb{Z}[1, 50]$. For each training image, the percentage of SIFT *features* matched to each visual word is calculated. This information can be represented by a histogram of size 1×50 (Figure 7.16g). These histograms are referred to as fingerprints. Ideally, melt pools with similar in-situ appearances will have similar *fingerprints*, while melt pools with dissimilar appearances will have dissimilar *fingerprints*. The observant reader may have noticed that in addition to the loss of scale information (desired), the method presented above also loses information about the relative spatial positions of the SIFT *features* (not desired). To mitigate this side-effect, the spatial relationships are partially represented by a multi-modality histogram which is detailed in the following subsection. While not explored in this work, another method for representing *visual word* occurrences known as VLAD (Vector of Locally Aggregated Descriptors) [217] may also be effective for this application.

7.3.5 Multi-Modality Representation of Melt Pool Morphology

While each melt pool could potentially be represented by a single histogram of *visual word* occurrences (as is the case for each *patch* in Section 4.3), such a representation was found to provide relatively poor differentiation between certain in-situ melt pool appearances. To improve differentiation, additional information contained within the original melt pool images was incorporated into the melt pool morphology representation by appending multiple histograms of *visual word* occurrences (Figure 7.16g); where the calculation of each histogram was preceded by a different set of pre-processing operations performed on the original (coaxially-transformed) melt pool image.

As noted in Section 7.2.2, the information of interest in each melt pool image spans a wide dynamic range. In order to capture gradient fields across this dynamic range, three different contrast adjustments [156] are applied to each melt pool image in the training database. Specifically, the pixel-wise data in each image are scaled using gamma values¹²⁷ of 1 (no change), 0.75 (decreased contrast), and 10 (increased contrast). Physically, decreasing the contrast allows gradient fields to be captured for the lower temperature (strictly, lower emitting) regions of the image while increasing the contrast emphasizes the gradient fields in the higher temperature (strictly, higher emitting) regions of the image. In other words, the diffuse vapor plume and colder spatter particles may be more visible in the lower contrast image, while only the melt pool body and the hottest spatter particles will be visible in the high contrast image. Note that this difference in "visibility" is also dependent upon the gradient magnitude threshold used to define a "strong edge" (see Section 7.3.3).

Recall that the *fingerprint* presented in the previous subsection contains no information about the relative spatial configuration of the SIFT *features*. In order to preserve some of this spatial information, each of the three contrast-adjusted melt pool images is segmented into three different components. First, the spatter is isolated using the same connectedcomponents algorithm described in Section 7.2.4. Then the main melt pool body (i.e. everything that is not considered spatter) is separated into the "tail" region and the "nose" region. These two regions are delineated by the line perpendicular to the ξ -axis at the point of maximum melt pool width. This process, shown schematically in Figure 7.19, produces nine distinct

¹²⁷ The pixel-wise scaling is accomplished using a non-linear function of the form: $I_{out} = I_{in}^{\Gamma}$, where I_{in} is the original value of the pixel, I_{out} is the adjusted value of the pixel, and the shape of the curve is defined by Γ [156].

fingerprints (calculated as described in Section 7.3.4) for each training image. All nine *fingerprints* are combined (appended) to form the final multi-modality representation of the melt pool which is of size 1×450 . Note that programmatically all of the SIFT *features* are first calculated across the entirety of each contrast-adjusted melt pool image before segmentation in order to avoid the creation of artifacts at the boundaries of the segmented images.



Figure 7.19: False-color and coaxially-transformed images of a melt pool produced using PV #1. The top row shows the melt pool image after three different contrast adjustments. The middle row shows each contrast-adjusted image segmented into spatter, the nose region, and the tail region. The bottom row shows a visualization of the SIFT *features* generated using a MATLAB script written by Dollar [214].

While the choice of the three segmentation regions (spatter, nose, and tail) is informed by

knowledge of melt pool dynamics (e.g. that spatter is likely to occur when a deep keyhole is

present [121] and balling affects the morphology of the tail [88], [107]), the choices of the exact

contrast adjustments¹²⁸, SIFT window size, and number of *visual words*, are less informed. Therefore the author suspects that an improved representation of the in-situ melt pool appearance could be developed through the use of Deep Learning techniques and this topic is discussed further in Section 8.3.

7.3.6 Training

Unlike in Chapter 4, the ground-truths in this work are not known a priori. That is, the author does not know what groupings are appropriate to describe the melt pool morphologies; much less which specific in-situ melt pool appearances should correspond to each label. When confronting an "unsupervised" ML task, it is often useful to visualize the locations of the final *feature* vectors (in this case the *fingerprints*) of each training datum in *feature* space. Because *feature* space is typically high-dimensional (in this case 450D) direct visualization is not possible and a low-dimensional approximation of *feature* space must be used instead. A common algorithm for constructing this low-dimensional approximation is t-SNE (t-distributed Stochastic Neighbor Embedding) [218]. In a low-dimensional t-SNE visualization the relative distances

¹²⁸ An informed choice of the low contrast adjustment value was attempted by converting the pixel intensity values to emissive temperatures using the inverse Sakuma-Hattori equation [232] as described by Fisher [34]. However this adjustment was found to be detrimental to differentiation between the in-situ appearances of the melt pools and was not included in the final multi-modality *fingerprint*.

between the *feature* vectors are preserved¹²⁹ (albeit non-linearly), however the absolute distances between data and their relative positions are lost [218]. In other words, *fingerprints* that are close in 450D space will be close in t-SNE 2D space and *fingerprints* that are far apart in 450D space will be far apart in t-SNE 2D space, but no other conclusions about their relative distribution can be inferred. For this reason, t-SNE visualizations are often used to identify natural clusters of high-dimensional *feature* vectors [218]. Figure 7.20 shows a t-SNE visualization of the 24,385 training *fingerprints* in which each datum has been color-coded using the only label available – the dominant ex-situ morphological characteristics for the given parameter combination (see Section 6.3.5).

¹²⁹ Preservation of the relative distances between high-dimensional data is accomplished using statistical operations [218]. Conceptually, consider the following: You are "standing" at a datum in high-dimensional space and you wish to reach out your hand and grab a different datum. Now, if you describe your ability to grab a specific datum using a high-dimensional Gaussian function, then you will be more likely to grab a datum near you than a datum farther away from you. Therefore, in order to map the relative distances between data to low-dimensional space this behavior should be similar in both representations. That is, if you stand on the same datum in low-dimensional space and reach out your Gaussian-determined hand, the probability of you grabbing each of the surrounding data should be similar to the probabilities found in the original high-dimensional space for those same surrounding data.



Figure 7.20: A t-SNE visualization of all of the training *fingerprints* with each datum color-coded according to the dominate ex-situ morphology of melt pools produced using the same process parameter combination. While t-SNE visualizations are sometimes presented with axes, they have been removed from this figure as neither their magnitudes nor their relative values carry any physical meaning. The t-SNE algorithm was executed with a perplexity of 75 while all other parameters were set to their default value [218] and no dimensionality reduction using PCA was implemented.

Perhaps unsurprisingly, the above t-SNE visualization is not particularly illuminating. For example, *fingerprints* of melt pools produced using *desirable* process parameters overlap extensively with *fingerprints* of melt pools produced using *balling* and *keyholing porosity* process parameters. As discussed in Chapter 6, many of the ex-situ flaws such as keyholing porosity and balling are periodic in nature, therefore it is to be expected that each process parameter combination will produce a range of in-situ melt pool morphologies.

A far more effective approach is to visualize the regions in process space for which a given set of *fingerprints* appear. In order to accomplish this, "sets" of similar *fingerprints* must first be identified and delineated. Grouping of similar *fingerprints* was performed using a standard kmeans unsupervised clustering algorithm [163], represented by Figure 7.16h. For this work, cluster initialization was performed using random seeding, with preference given to a uniform spacing between clusters. Ideally, the number of requested clusters could be informed by the t-SNE visualization in Figure 7.20, unfortunately, very few distinct clusters are evident. Instead, many of the melt pool *fingerprints* appear to exist on a continuum. Therefore, a total of 30 clusters were requested; while this value is somewhat arbitrary, poor differentiation between in-situ melt pool appearances was observed when fewer than 20 clusters were delineated while additional meaningful differentiations did not appear when more than approximately 30 clusters were delineated. Cluster seeding is repeated 10 times to reduce the chance of the algorithm converging to a poor solution; e.g. a shallow local minimum instead of a global, or at least a deeper local, minimum.

Figure 7.21 shows the percentage of E2 melt pools at each of the 29 process parameter combinations with *fingerprints* belonging to one of the 30 clusters. Observe that this set (cluster) of *fingerprints* is far more prevalent in the *keyholing porosity* regime than elsewhere in process space. Therefore this set of *fingerprints* (i.e. this in-situ melt pool appearance) can be linked to *keyholing porosity*. This linkage process (Figure 7.16i) is performed by a human and is repeated for all 30 sets (clusters) of *fingerprints*. Over the course of this work it was observed that no clusters of *fingerprints* could be associated with only the *severe keyholing* region of process space, implying that the depth of the keyhole-mode melting vapor cavity does not have a controlling influence on the in-situ appearance of the melt pool. Furthermore, it was observed that a side-effect of the multi-modality *fingerprint* histogram is the occurrence of clusters of *fingerprints* with locations in high-dimensional space that are driven primarily by the morphology of the spatter, which is not necessarily correlated with any ex-situ flaws (Section

258

6.3.5). Clusters of this type were labeled *spatter* are discussed in more detail in Section 7.4.1. The implications of this side-effect and possible solutions are discussed in Section 8.3.





prevalence of a given set (cluster) of melt pool fingerprints across process space. Values within the text boxes indicate the percentage of melt pools imaged during E2 (at each process combination) parameter with fingerprints belonging to the given cluster. The ex-situ morphology information from Chapter 6 is overlaid on top of the heat map. In-situ data were not collected for the seven parameter combinations (circles) without an associated text box (see Section 7.2.6).

Once unsupervised ML techniques have been used to understand and label the in-situ appearances of the E2 melt pools, supervised ML techniques can be employed to classify melt pools not included in the training database. It is extremely important to note that any classifications based on this approach are not tied directly to ex-situ outcomes. In other words, even if a melt pool is classified as *keyholing porosity* it cannot be concluded that keyholing porosity was indeed generated by that melt pool – even within some degree of uncertainty. Instead, such a classification only indicates that a given melt pool has an in-situ *fingerprint* which is similar to *fingerprints* found most prevalently (or ideally, exclusively) in the *keyholing porosity* regime of process space.

Classification of the melt pool morphologies (i.e. the 450 element *fingerprints*) is performed using a multi-class Support Vector Machine (SVM) [219], [220] which is shown

graphically in Figure 7.16j. In its original formulation, an SVM is a binary classifier capable of distinguishing between only two different classes [220]. During training an SVM learns a hyperplane which bisects the high-dimensional *feature* space such that all of the *feature* vectors belonging to one class lie on a different side of the hyperplane than all of the *feature* vectors belonging to the other class (with the minimum error possible) [220]. A variety of methods are available to apply SVMs to non-binary (i.e. multi-class) classification problems. Perhaps the simplest such method converts a multi-class problem into a set of binary classification problems such that binary classifiers can be trained to distinguish between each class and "all of the other classes" [221]. The multiple binary classifiers can then be combined to form the multi-class SVM. All of the training parameters for the multi-class SVM were set to the default values listed in [219].

7.3.7 Melt Pool Classification and Performance of the ML Methodology

Once the training of the ML methodology is complete, melt pool images captured during E3a, E3b, and E3c can be classified as one of five melt pool types: *desirable, balling, undermelting, keyholing porosity,* or *spatter* (Figure 7.16k). Each E3 melt pool image is first warped and transformed as described in Sections 7.2.3 and 7.2.4. Note that the appropriate in-plane rotation angle was determined a priori to improve the robustness of the transformation operation. After the coaxial transformation, the multi-modality *fingerprint* is calculated (Sections 7.3.3 – 7.3.5) and classified by the trained SVM (Section 7.3.6). The results of these classifications are presented in Sections 7.4.2 – 7.4.4. The classification process requires approximately 0.4 seconds per frame on a single Intel[®] i7-6700K 4.00 GHz processor. It is important to note that in the absence of ground-truth labels, quantifying the overall performance of the presented ML methodology is not possible. However, 10-fold cross-validation [179], [180, p. 78] was performed during training of the multi-class SVM and reported a classification accuracy of 85.1%. In other words, the hyperplanes learned by the SVM are able to properly delineate the *fingerprints* contained within the E2 training database according to the labels applied in Section 7.3.6 with an accuracy of 85.1%. Additional details regarding k-fold cross-validation and other ML performance metrics are available in Section 4.5.

7.4 Results

7.4.1 Melt Pool Morphologies across Process Space (Experiment 2)

At this point the reader is encouraged to review the ex-situ melt pool morphologies discussed in Section 6.3.5. The prevalence of a particular in-situ melt pool appearance throughout process space can be visualized by considering the cluster assignments at each tested process parameter combination. In other words, for each parameter combination imaged during E2, a certain percentage of the captured melt pool frames will belong to a given set (cluster) of *fingerprints* described in Section 7.3.6. For the remainder of this chapter, a set of similar *fingerprints* is referred to as a *morphology*. While the example heat map shown in Section 7.3.6 (Figure 7.21) contains information about only one of the *morphologies*, Figures 7.22, 7.24, 7.26, 7.28, and 7.30 show the percentage of E2 melt pools (at each process parameter combination) belonging to the *morphologies* associated with a given ex-situ

morphology. Also as in Section 7.3.6, the ex-situ morphology results are overlaid on top the heat map.

Figure 7.22, for example, shows the distribution of the six melt pool *morphologies* (i.e. six sets of similar *fingerprints*) associated with *desirable* outcomes. Observe that the *desirable morphologies* are most prevalent in the "center" of studied processing space – away from the high energy density, low energy density, and balling regimes. The occurrence of these *desirable morphologies* ranges from 21% at PV #19 (250 W, 200 mm/s) to 77% at PV #22 (250 W, 800 mm/s). During E2, 76% of the melt pools produced using the EOS nominal PV #1 parameters (285 W, 960 mm/s) had *fingerprints* associated with *desirable* outcomes. More strictly, 76% of these melt pools had *fingerprints* belonging to one of the six clusters associated with *desirable* outcomes. Figure 7.22 shows several examples of *desirable* melt pools.





Figure 7.22: A heat map showing the prevalence of melt pool morphologies associated with *desirable* ex-situ outcomes. Values within the text boxes indicate the percentage of melt pools imaged during E2 (at each process parameter combination) with desirable morphologies. The ex-situ morphology information from Chapter 6 is overlaid on top of the heat map. In-situ data were not collected for the seven parameter combinations (circles) without an associated text box (see Section 7.2.6).



Figure 7.23: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *desirable* ex-situ outcomes.

Figure 7.24 shows the distribution of the single melt pool *morphology* associated with *balling*. Observe that the *balling morphology* is most prevalent in the high beam power and high beam velocity regime of process space. The occurrence of this *balling morphology* ranges from 2% for PV #7 (150 W, 200 mm/s) to 26% at PV #36 (370 W, 1400 mm/s). During E2, 19% of the melt pools produced using the EOS nominal PV #1 parameters (285 W, 960 mm/s) had *fingerprints* associated with *balling*. Figure 7.25 shows several examples of *balling* melt pools. In general, the *balling* melt pools are more elongated than the *desirable* melt pools (this difference is expected as discussed in Section 6.1). In some of the high speed images (Figure 7.25) the balling instability itself is visible as small circle separated from the main melt pool body and located directly behind (negative ξ -axis) the melt pool tail.



Figure 7.25: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *balling*. Note the balling instabilities visible just behind the melt pool tail in the left and center images.

Figure 7.26 shows the distribution of the two melt pool *morphologies* associated with *under-melting*. Observe that the *under-melting morphologies* are most prevalent in the low beam power regime of process space. Interestingly, these *morphologies* extend well into the regions of process space producing melt pools with depths greater than the 70 µm effective powder layer thickness. As noted in Section 7.2.6, in-situ data were not successfully collected

for much of the *under-melting* region; more in-situ data may allow for a more robust understanding of morphologies associated with *under-melting*. The occurrence of these *undermelting morphologies* ranges from 0% at many parameter combinations to 3% at PV #4 (100 W, 600 mm/s). During E2, 1% of the melt pools produced using the EOS nominal PV #1 parameters (285 W, 960 mm/s) had *fingerprints* associated with *under-melting*. Figure 7.27 shows several examples of *under-melting* melt pools. In general, *under-melting* melt pools exhibit a "fragmented" appearance composed of multiple irregular bodies of comparatively similar size.



EOS Nominal Parameter Combination
Desirable
Balling
Under-Melting
Severe Keyholing
Keyholing Porosity

Figure 7.26: A heat map showing the prevalence of melt pool morphologies associated with under-melting. Values within the text boxes indicate the percentage of melt pools imaged during E2 (at each process parameter combination) with under-melting morphologies. The ex-situ morphology information from Chapter 6 is overlaid on top of the heat map. In-situ data were not collected for the seven parameter combinations (circles) without an associated text box (see Section 7.2.6).



Figure 7.27: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *under-melting*.

Figure 7.28 shows the distribution of the three melt pool *morphologies* associated with *keyholing porosity*. Observe that the *keyholing porosity morphologies* are most prevalent in the low beam velocity regime of process space – particularly at higher beam powers. Refer to Section 6.3.5 for a discussion of keyholing porosity versus keyhole-mode melting. The occurrence of these *keyholing porosity morphologies* ranges from 0% at PV #36 (370 W, 1400 mm/s) to 66% at PV #19 (250 W, 200 mm/s). During E2, 0% of the melt pools produced using the EOS nominal PV #1 parameters (285 W, 960 mm/s) had *fingerprints* associated with *keyholing porosity*. Interestingly, no in-situ *morphologies* could be associated with the *severe keyholing* melt pools discussed in Section 6.3.5. This suggests that the *keyholing porosity morphologies* are detectable due to the instability, and not the depth, of the vapor cavity present during keyhole-mode melting [85]. Figure 7.29 shows several examples of *keyholing porosity* melt pools. In general, *keyholing porosity* melt pools.



Figure 7.29: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *keyholing porosity*.

Figure 7.30 shows the distribution of the eighteen melt pool *morphologies* associated with *spatter*. Strictly, *spatter* is not a melt pool morphology identified by the ex-situ analysis presented in Chapter 6. Indeed, its appearance as unique in-situ *morphologies* is an artifact of the multi-modality melt pool representation discussed in Section 7.3.5. While independent calculation of a spatter *fingerprint* was observed to improve overall classification performance, it occasionally allows the morphology of the spatter to "overpower" information about the

shape of the melt pool itself – resulting in the aforementioned *spatter morphologies*. Potential approaches for addressing this challenge are discussed in Section 8.3. In general, *spatter* appears least frequently in the region of process space associated with *desirable* melt pool morphologies. The occurrence of these *spatter morphologies* ranges from 3% at PV #33 (370 W, 800 mm/s) to 12% at several parameter combinations. During E2, 4% of the melt pools produced using the EOS nominal PV #1 parameters (285 W, 960 mm/s) had *fingerprints* considered to be driven by their spatter morphology. Figure 7.31 shows several examples *spatter* melt pools. Observe that the melt pool bodies exhibit dramatically different melt pool morphologies while the primary commonality between the selected images is the presence of spatter.





Figure 7.30: A heat map showing the prevalence of melt pool *morphologies* associated with *spatter*. Values within the text boxes indicate the percentage of melt pools imaged during E2 (at each process parameter combination) with *spatter morphologies*. The ex-situ morphology information from Chapter 6 is overlaid on top of the heat map. In-situ data were not collected for the seven parameter combinations (circles) without an associated text box (see Section 7.2.6).

^Dercentage of Melt Pools



Figure 7.31: A selection of false-color melt pool images captured by the high speed camera during E2. All three melt pools have *fingerprints* associated with *spatter*.

The melt pool *fingerprints* themselves (i.e. the histograms of size 1×450) can be plotted in low-dimensional space using the t-SNE algorithm (first introduced in Section 7.3.6). While the fingerprints shown in Figure 7.20 are color-coded based the dominant ex-situ morphological characteristics for the given parameter combination, the *fingerprints* in Figure 7.32 are instead color-coded based on their associated in-situ melt pool morphology. Because the fingerprints have been color-coded based on the cluster to which they belong, it is wholly unsurprising that groupings of color are readily apparent in the figure below. Indeed, the reader should be careful to not draw any conclusions based solely on the fact that clusters are extant. Nonetheless, this t-SNE representation is illuminating on several fronts. For example, while multiple distinct morphologies (clusters) were found to be indicative of specific types of melt pools (i.e. three clusters are associated with keyholing porosity melt pools), fingerprints belonging to those morphologies tend to lie close to fingerprints belonging other morphologies associated with that ex-situ melt pool type. For example, the several keyholing porosity morphologies are closer to each other than to the balling morphologies, as evidenced by the relatively contiguous grouping of red-colored *fingerprints*. Conversely, not only are the unique

spatter morphologies numerous and highly distinct from the other *morphologies*, they are also quite different from each other. This is not surprising given the origin of the *spatter morphologies* which is discussed earlier in this subsection.



Figure 7.32: A t-SNE visualization of all of the training *fingerprints* with each datum color-coded according its in-situ morphology. While t-SNE visualizations are sometimes presented with axes, they have been removed from this figure as neither their magnitudes nor their relative values carry any physical meaning. The t-SNE algorithm was executed with a perplexity of 75 while all other parameters were set to their default value [218] and no dimensionality reduction using PCA was implemented.

Finally, Figure 7.33 shows the occurrence of the various melt pool morphologies for several

selected process parameter combinations. As hypothesized and discussed in Chapter 6, every process parameter combination produces melt pools with a range of in-situ appearances. For example, even the parameter combination producing the most *keyholing porosity* melt pools (PV #19) also produces melt pools with *desirable* morphologies in 21% of frames. This behavior further illustrates the difficulties associated with directly using the annotated t-SNE plot in Figure 7.20 to train the melt pool classification algorithm.



Figure 7.33: The melt pool morphologies produced by four different process parmeter combinations during E2. PV #1 is the EOS nominal parameter combination. PV #36 produced the largest percentage of *balling* melt pools. PV #4 produced the largest percentage of *under-melting* melt pools. PV #19 produced the largest percentage of *keyholing porosity* melt pools.

7.4.2 Stripe Edges (Experiment 3a)

Figure 7.33 reports the melt pool morphology classifications near the edge of a 10 mm wide stripe. As discussed in Section 7.2.7, during E3a the FoV was centered on the stripe edge, allowing data collection over a beam travel distance of approximately 3.5 mm. Data were collected from multiple adjacent melt tracks as the laser beam propagated along the stripe direction (Figure 1.8). Each data bar reports the classifications for all of the melt pools within a certain distance of the stripe edge (indicated by the vertical dashed line); data are binned in increments of approximately 500 µm. Data were collected for melt pools traveling both toward and away from the stripe edge with each classification average based on between 19 and 45 frames of data. No significant morphological differences are apparent between melt pools 3.0

mm to 3.5 mm and 0.0 mm to 0.5 mm away from the edge. Notably, however, the E3a melt pools were much more likely to be classified as *keyholing porosity* and much less likely to be classified as *balling* than the E2 melt pools produced using the same EOS nominal process parameter combination (PV #1) (see Figure 7.33). Therefore it is possible that the stripe edge influences melt pool morphology at a distance exceeding 3.5 mm, and therefore the transition of interest occurred outside of the high speed camera's FoV.



Figure 7.34: Melt pool morphology classifications near the edge of a stripe that is 10 mm wide. Each data bar bins classifications in 500 μ m increments from the stripe edge (indicated by the vertical dashed line). The values in parantheses indicate the number of melt pools classifications included in the corresponding bin. The indicated positive *y*-axis corresponds to the global coordinate system used throughout this thesis.

Figure 7.35 reports the melt pool morphology classifications near the edges of stripes ranging in width from 10 mm (EOS nominal) to a width of only 0.5 mm. Data were collected for melt pools traveling both toward and away from the stripe edges. Note that fewer data are available as the stripe width decreases and that the upper bar plot is a simplified duplicate of the plot shown above in Figure 7.34. No significant morphological differences are apparent between melt pools rastering within the 10 mm stripe and melt pools rastering within the 0.5 mm stripe. Just as observed for the 10 mm wide stripe, the melt pools were much more likely to be classified as *keyholing porosity* and much less likely to be classified as *balling* than the E2 melt pools produced using the same EOS nominal process parameter combination (PV #1) (see Figure 7.33).



Figure 7.35: Melt pool morphology classifications near the edge of stripes that are between 10 mm and 0.5 mm in width. Each data bar bins classifications in 500 μ m increments from the stripe edge (indicated by the vertical dashed lines). Note that the annotated locations of the stripe edges are approximate and may be shifted slightly relative to the data bars. This approximation may be partially responsible for the appearance of melt pools existing beyond the nominal stripe boundaries in the case of the 0.5 mm wide stripe. The values in parantheses indicate the number of melt pools classifications included in the corresponding bin. The indicated positive *y*-axis corresponds to the global coordinate system used throughout this thesis.

Finally, Figure 7.36 reports the melt pool morphology classifications at several heights during the printing of a thin wall structure that is 0.5 mm in width. Note that the data collected at the thin wall height of 0 mm are the same data reported above in Figure 7.35 for the 0.5 mm wide stripe. Data were collected for melt pools traveling in both directions with each classification average based on between 31 and 37 frames of data. No significant morphological differences are apparent between the melt pools used to fuse the first layer of the thin wall and the melt pools used to fuse the final layer of the thin wall – 5 mm above the bulk E3a test

artifact. As in the previous E3a results, the melt pool morphologies differed from those observed during E2 for the same EOS nominal process parameter combination (PV #1) (see Figure 7.33).



Figure 7.36: Melt pool morphology classifications during the construction of a thinwall structure that is 0.5 mm in width. Each data bar reports classifications at heights ranginging from 0 mm to 5 mm above the bulk E3a test artifact. The values in parantheses indicate the number of melt pools classified for the corresponding build layer. The indicated positive *z*-axis corresponds to the global coordinate system used throughout this thesis.

Whether the observed differences between the E3a morphologies and the equivalent E3 morphologies ae due to the proximity of the melt pools to the stripe edge (within 3.5 mm) or the minimal spatial and temporal separation between adjacent melt tracks compared to those in E2 (110 μ m versus 500 μ m and a minimum 140 ms delay) is unknown at this time. It is worth noting that, at the current camera frame rate, the nominal melt pool travels 150 μ m between each frame. Additionally, it is possible that significant morphological changes triggered by proximity to the edge of the stripe may occur over distances comparable to, or less than 150 μ m. Therefore it may be advantageous to study this particular non-bulk geometry using a

higher camera frame rate. Furthermore, the nominal parameter combination's distance from the keyholing porosity regime (Figure 6.28) suggests that it may not be an ideal candidate¹³⁰ for observing morphological changes due to local preheating from adjacent melt tracks [184].

7.4.3 Overhang Region (Experiment 3b)

Figure 7.37 reports the melt pool morphology classifications during fusion of the first layer of a 5.0 mm wide unsupported overhang. Recall that in the first layer of an overhang there is no fused material below the melt pool, only unfused powder. Data were collected from multiple adjacent melt tracks as the layer was fused. Each data bar reports the classifications for all of the melt pools that are certain distance across the width of the overhang; data are binned in increments of approximately 500 µm. The edges of the overhang are indicated by the vertical dashed lines. Note that the left-most and right-most data bars report data corresponding to melt pools beyond the extents of the overhang, i.e. those melt pools are, at least partially, on top of fused material. Data were collected for melt pools traveling in both directions across the overhang with each classification average based on between 60 and 129 frames of data. It is readily apparent that the melt pools transition from primarily *desirable* to primarily *keyholing* porosity as the melt pool travels from the edge of the overhang to its center. A symmetric transition back to *desirable* classifications occurs as the melt pool exits the overhang region and returns to previously-fused material. This behavior is fully expected as the low thermal conductivity [100], [210] of the unfused powder is expected to result in a much deeper and

¹³⁰ It was not known a priori that the nominal parameters would be so far removed from the keyholing porosity regime. The prior work in the literature [12, Ch. 6] focused primarily on delineating the keyhole-mode melting regime (without the presence of a powder layer) and did not explore stable, non-porosity induced keyholing-mode melting.

more unstable keyhole vapor pocket than is otherwise produced by the EOS nominal process parameter combination (PV #1).



Figure 7.37: Melt pool morphology classifications across the 5.0 mm width of the first layer of an unsupported overhang. Each data bar bins classifications in 500 μ m increments along the melt tracks spanning the width of the overhang. The extents of the overhang are indicated by the vertical dashed lines. The values in parantheses indicate the number of melt pool classifications included in the corresponding bin. The indicated positive *x*-axis corresponds to the global coordinate system used throughout this thesis.

Figure 7.38 reports the melt pool morphology classifications during fusion of the first five layers of a 5.0 mm wide unsupported overhang. For the first three to four layers, the observed melt pool behavior is similar to that discussed above. That is, a substantial percentage of the melt pools transition from *desirable* to *keyholing porosity* as they pass over the unsupported overhang. Note that that the upper bar plot is a simplified duplicate of the plot shown above in Figure 7.37. As expected, the melt pools spanning the overhang are increasingly classified as *desirable* as subsequent layers are built. In other words, the morphology of the melt pools improves as the E3b test artifact approaches bulk geometry. Indeed, by the fifth layer melt pool behavior in the center of the overhang is quite similar to melt pool behavior near the edges.

Interestingly, the melt pools fusing the fifth layer were approximately 160 μ m above the powder bed and the average melt pool depth for the EOS nominal process parameter combination (PV #1) was measured to be 150 μ m. This suggests that the thermal influence of an overhang may extend a vertical distance comparable to the nominal depth of the melt pool.



Figure 7.38: Melt pool morphology classifications across the 5.0 mm width of the first five layers of an unsupported overhang. Each data bar bins classifications in 500 μ m increments along the melt tracks spanning the width of the overhang. The extents of the overhang are indicated by the vertical dashed lines. The indicated positive *x*-axis corresponds to the global coordinate system used throughout this thesis.

7.4.4 Contours (Experiment 3c)

Figure 7.39 reports the melt pool morphology classifications for the exposure of three contour passes using the EOS nominal bulk process parameter combination (PV #1). The contour passes ranged from 0 μ m to 150 μ m from the nominal edge of the E3c test artifact. Each data bar reports the classification of 65 frames of data captured over two sets of contour exposures as described in Section 7.2.7. No significant morphological differences are apparent between melt pools on the edge of the E3c test artifact versus melt pools 150 μ m from the

edge. Furthermore, the morphology classifications are similar to those observed for the same parameter combination (PV #1) during E2 (see Figure 7.33). Such a lack of differentiation is perhaps surprising given the differing thermal conditions near the free edge (due to the low thermal conductivity of the powder [100], [210]). However, as noted previously, the EOS nominal parameter combination is well distant from the keyholing porosity regime and therefore it may not be an ideal candidate¹³¹ for observing morphological changes induced by a melt pool's proximity to a free edge.

Figure 7.40 reports melt pool morphology classifications during the exposure of three subsequent contour passes using a set of process parameter combinations designed by the author and Dr. Sneha Prabha Narra of CMU to mitigate near-surface keyholing porosity. The contour passes ranged from 0 µm to 140 µm from the nominal edge of the E3c test artifact. Each data bar reports the classification of between 249 and 252 frames of data captured over two sets of contour exposures as described in Section 7.2.7. Despite the stated design intention, a substantial percentage of the melt pools are indeed classified as *keyholing porosity*. Upon further review, it was determined that the designed process parameters lie very near to the keyholing porosity regime reported in Figure 6.28. Furthermore, it is considered likely that the proximity of the melt pools to a free edge increased the depth of the keyholing vapor pocket, further exacerbating the problem. The design of these process parameters was originally informed by the process map reported by Narra [12, Ch. 6], based on which they were

¹³¹ It was not known a priori that the nominal parameters would be so far removed from the keyholing porosity regime. The prior work in the literature [12, Ch. 6] focused primarily on delineating the keyhole-mode melting regime (without the presence of a powder layer) and did not explore stable, non-porosity induced keyholing-mode melting.

expected to be well removed from the keyholing regime. These results further motivate the need for careful study of the effects of powder on melt pool size and morphology (Section 6.3.4) as well as the necessity of in-situ monitoring of the process.



Figure 7.39: Melt pool morphology classifications for three contour passes ranging from 0 μ m to 150 μ m from the free edge. All of the contours were exposed using PV #1. The values in parantheses indicate the number of melt pool classifications for the corresponding contour pass. The process parameter combinations are also noted.



7.5 Proposed Real-Time Implementation Strategies

While the work presented heretofore in this chapter provides valuable insight into melt pool-scale flaw formation mechanisms in L-PBF, it is not immediately deployable as a component in a real-time closed-loop control system. The extremely high data rates and the computational burden of the presented methodology preclude such a direct implementation. Indeed, the 0.4 seconds required to classify a melt pool (Section 7.3.7) is three to four orders of magnitude too slow for a real-time application. However, it is the author's opinion that much of the algorithm could conceivably be implemented in hardware (i.e. circuitry) which would dramatically reduce the computation time required to convert a melt pool image (on the order of 10⁴ to 10⁶ integers) to a *fingerprint* describing its morphology. This single vector could then be used in a control loop even at high sampling rates (on the order of 10 kS/s to 100 kS/s). In fact, the real-time calculation of L-PBF in-situ melt pool size has already been demonstrated by Clijsters et al. [81] and similar systems, combined with high speed pyrometer data (also a single number per sample), are currently being used by several L-PBF machine manufacturers for process control [145], [146]. While a complete hardware implementation of the proposed methodology is well beyond the scope of this thesis, a potential FPGA (Field Programmable Gate Array) [222] architecture is outlined below.

Figure 7.41 provides a high-level schematic of a potential FPGA implementation. Pixel-wise data are pulled directly off the camera sensor and are buffered by a latch. When triggered, the output pins of the latch are updated and read by the FPGA (Figure 7.41a). At this point, the horizontal and vertical gradients at each pixel are calculated simultaneously; that is, if there are 10^6 pixels, 2×10^6 gradient calculations will be performed in parallel. The pixel-wise gradient magnitudes and orientations are also calculated in parallel. The pixel-wise orientations at strong gradients are binned as described in Section 7.3.3. Next (Figure 7.41b), the SIFT descriptors are calculated and converted into HOG vectors in a window-wise fashion; that is, if each window is of size 2 pixels \times 2 pixels, 2.5×10^5 HOG vectors will be constructed simultaneously. The error between each HOG vector and each *visual word* in the *dictionary* is

280

then calculated and the closest *visual word* is identified; the author expects that this operation will be the most computationally expensive of all of the proposed parallelized calculations. After the nearest *visual word* is assigned to each SIFT window (Figure 7.41c), the occurrence of each *visual word* in the melt pool image is counted and stored in the final *fingerprint* describing the melt pool morphology. This *fingerprint* is then sent from the FPGA to an SVM classifier implemented on a standard CPU (Figure 7.41d). Observe that if the *fingerprint* is of size 10³, a 10³ fold reduction in data transmission burden has been achieved versus the direct transmission of a raw of image of size 10⁶ pixels. Finally, the classification of the melt pool morphology may then be used to inform a feedback controller in order to flag a potential defect or even trigger an autonomous "re-work" [133] strategy.






Figure 7.41: A high-level schematic for a potential FPGA implementation of the presented melt pool morphology classification methodology.

Also of note, work is ongoing at CMU by the author, Dr. Zhou Chen, Prof. Mahadev Satyanarayanan, and Prof. Jack Beuth to use "cloudlet computing" to enable real-time classification of melt pools at a reduced frame rate. In other words, while the high speed camera's exposure time would remain low, the effective classification rate would be on the order of only 10 fps. Such a system may allow for useful information to be gleaned, in real-time, about melt pool behavior during a build without the development of an FPGA architecture.

7.6 Discussion and Summary

In this chapter, In718 L-PBF melt pools were imaged using a high speed, visible-light camera with a fixed Field of View. Analysis of the melt pool images required the development of a custom algorithm to transform the imaged melt pools from an Eulerian reference frame to a Lagrangian (coaxial) reference frame. High speed camera data collected from across process space were analyzed using contemporary Computer Vision *feature* extraction methods and unsupervised Machine Learning techniques. These data were then used to link in-situ melt pool morphologies to the ex-situ melt pool morphologies and processing defects characterized in Chapter 6. These linkages allowed for the identification of potential in-situ flaw formation signatures and the training of a melt pool classification algorithm.

Feature extraction was performed using SIFT and the BoW ML technique – resulting in a scale-invariant representation of the in-situ melt pool morphology. The need for scale-invariance is well demonstrated by Figure 7.23, in which melt pools of dissimilar size are linked to similar ex-situ outcomes. In order to compensate for the loss of spatial information during *feature* extraction as well as the large dynamic range of the input images, the final *fingerprint* describing melt pool morphology is composed of nine different segments. Each segment describes the gradient fields under different contrast adjustments and in different regions of the melt pool. While effective at differentiating between melt pools of differing in-situ appearances, many of the parameters controlling the architecture of the final morphology

descriptor remain un-optimized; Deep Learning techniques (such as those described in Section 4.4) may be well suited for developing an even more robust representation of in-situ melt pool morphology.

Several in-situ melt pool appearances were successfully linked to the ex-situ melt pool types identified in Chapter 6: (1) A set of in-situ morphologies were found to occur almost exclusively at keyholing porosity parameter combinations (Figure 7.28). (2) Conversely, no insitu morphologies were found to be associated with severe keyholing. (3) A set of in-situ morphologies were found to occur more frequently in the balling region of process space (Figure 7.24). (4) Interestingly, the set of in-situ morphologies most closely linked to undermelting extended well throughout the low power regime of process space (Figure 7.26). (5) Several melt pool morphologies were associated with desirable ex-situ outcomes (Figure 7.22). (6) As discussed in Section 7.3.6, an additional set of morphologies were associated with the presence of spatter, as opposed to a specific ex-situ outcome (Figure 7.30). The periodic nature of many of the flaws (Chapter 6) was observed in-situ, as each process parameter combination produced melt pools with *fingerprints* associated with a range of ex-situ morphologies. Using a unique approach, the associations determined via unsupervised learning were used as the inputs for a supervised learning technique; thereby enabling the classification of melt pools not studied during training.

Specifically, melt pools were imaged during the printing of three different test artifacts, each of which demonstrated a different non-bulk geometry. (1) The first test artifact (E3a) investigated the behavior of melt pools approaching a stripe boundary for stripes of varying widths. No significant morphological changes were observed either as the melt pools

approached the stripe edge (Figure 7.34) or as the stripe narrowed (Figure 7.35). Because morphological changes were anticipated, the author suggests that stripe boundaries be explored further using a variety of parameter combinations and a higher camera frame rate. (2) The second test artifact (E3b) investigated melt pools traveling across an unsupported overhang. Significant morphological changes were observed as the melt pools traveled away from the bulk region and across the overhang (Figure 7.37). After the overhang reached a thickness exceeding the nominal depth of the melt pool, the in-situ morphological behavior approached that observed in the bulk (Figure 7.38). (3) The third test artifact (E3c) investigated melt pool behavior during a contour pass. Morphological differences were not observed between the contours at the part edge and the contours 150 µm away from the edge (Figure 7.39); however this non-bulk geometry should be studied further with a variety of parameter combinations. Additionally, this analysis provides evidence that a specific set of process parameters originally designed to reduce near-surface porosity may still generate keyholing porosity (Figure 7.40); these results are also supported by the process map reported in Figure 6.28.

It is worth reiterating that any classifications based on the presented approach are not tied directly to ex-situ outcomes. In other words, even if a melt pool is classified as *keyholing porosity* it cannot be concluded that keyholing porosity was indeed generated by that melt pool – even within some degree of uncertainty. Instead, such a classification only indicates that a given melt pool has an in-situ *fingerprint* which is similar to *fingerprints* found most prevalently in the *keyholing porosity* regime of process space. Opportunities for improving the fidelity of

the classification methodology are discussed in Section 8.3. Finally, possible real-time implementations of the presented melt pool classification methodology are briefly discussed.

8 Conclusions

8.1 Summary

In order for Additive Manufacturing to successfully transition from a prototyping tool to a widely-deployed means of production, several obstacles must first be overcome. The Introduction of this thesis motivates two of these challenges: (1) increasing design freedom by expanding process space and (2) improving process stability, reducing operator burden, and ensuring part quality through the development of in-situ process monitoring methodologies. In this thesis, the expansion of Laser Powder Bed Fusion process space is studied via process mapping of the AlSi10Mg material system and the quantification of the effects of powder particle size on part quality and process stability. Contemporary Machine Learning and Computer Vision techniques are used to detect and classify L-PBF powder spreading anomalies – not only paving the way for real-time feedback control but also providing a powerful data analytics tool for studying the build process. Finally, process mapping of the Inconel 718 material system is combined with ML and CV analyses of visible-light high speed camera images in order to link in-situ melt pool morphologies with ex-situ process outcomes and study the fusion of non-bulk geometries.

In the second chapter, laser beam power and travel velocity are varied to develop process maps (with quantified levels of confidence) of cross-sectional melt pool geometry for the AlSi10Mg material system in an L-PBF process (Section 2.3.1). The variability and statistical behavior of melt pool geometry are studied across processing space and outliers are identified (Sections 2.3.2 and 2.3.3). Edge roughness is measured, although no trends were identified across process space (Section 2.3.5). Finally, bulk porosity arising from both the lack-of-fusion

and keyholing mechanisms is measured (Section 2.3.4) and combined with the aforementioned process maps in order to generate a robust processing window of process parameter combinations that are expected to produce parts with minimal porosity caused by either the lack-of-fusion or keyholing mechanisms (Section 2.3.6).

In the third chapter, the effects of three non-standard (i.e. not supplied by the machine manufacturer) Ti-6AI-4V powders and one non-standard Inconel 718 powder on L-PBF part and process quality are studied. Specifically, process parameters are modified in order to demonstrate successful manufacture of test specimens using powders with particle sizes up to 2.4 times larger than the largest particles found in the manufacturer-supplied Ti64 powder (Section 3.2.2). Part and build quality are both evaluated qualitatively. Powder spreading anomalies are quantified and a link between their spatial distributions and laser scan strategy is identified using the process monitoring approach presented in chapter four (Section 3.3.2). Bulk porosity produced via three different mechanisms is measured and four different edge roughness measures are calculated (Sections 3.3.3 and 3.3.4). Importantly, a correlation is identified between mean powder particle size and many of the reported build and part quality metrics (Section 3.3.5).

In the fourth chapter, autonomous powder bed anomaly detection and classification of several anomaly types is achieved through the use of Machine Learning and Computer Vision techniques applied to data collected by a low-resolution visible-light camera provided by the L-PBF machine manufacturer. Two different ML algorithms are applied to this problem: (1) the well-established Bag of Words approach (Section 4.3) and (2) a contemporary pre-trained Convolutional Neural Network (Section 4.4). Uniquely, the pre-trained CNN was modified

(MsCNN) in order to analyze data at multiple size scales, thereby dramatically improving anomaly classification accuracy. The performances of the two different ML methodologies are compared with the final MsCNN reporting an anomaly differentiation accuracy of 93% and a layer-wise analysis time of only 7 seconds (Sections 4.5.2 – 4.5.4). A single case study is presented in order to demonstrate the capabilities of the final methodology (Section 4.6).

In the fifth chapter, the powder bed monitoring work of chapter four is applied to ten different case studies in order to demonstrate the capabilities of the final methodology as well as to provide unique insights into the L-PBF building process. Specifically, thermal warping of small and large overhang structures is studied and linked to eventual *part damage* (Section 5.2.1). Delamination at the support-part interface (due to residual thermal stresses) is successfully detected even when the point of failure is well below the powder surface (Section 5.2.1). Build failure modes unique to high-aspect ratio and thin wall structures are identified (Sections 5.2.2 and 5.2.3). A correlation between detections of *super-elevation* and layer-wise energy density during fusion is observed for both the AlSi10Mg and In718 material systems (Section 5.2.4). *Super-elevation* is also linked to interactions between laser scan strategy and part geometry (Section 5.2.4). Finally, the general appearance of the powder bed across multiple material systems is explored (Section 5.2.5) and detection of an L-PBF machine malfunction via powder bed monitoring is demonstrated (Section 5.2.6).

In chapter six, a database of ex-situ melt pool morphology is developed for the Inconel 718 material system in an L-PBF process (Section 6.3.5). Unlike prior work in the literature, the studied melt pools were exposed in the presence of a powder layer. Special attention is given to the prevalence of processing flaws such as keyholing porosity and balling throughout process

space, as this information is critical for the development of the melt pool morphology classification algorithm presented in chapter seven. As in chapter two, the collection of data from multiple melt pool cross-sections allowed for the development of melt pool geometry process maps with quantified levels of confidence (Section 6.3.1) and the study of melt pool dimensional variability and statistical behavior (Sections 6.3.2 and 6.3.3). Finally, comparison of the developed process maps with In718 process maps reported in the literature suggests that the presence of a powder layer can have a significant effect on the cross-sectional depth of the melt pool (Section 6.3.4).

In chapter seven, classification of in-situ melt pool morphology (as captured with a high speed visible-light camera) was achieved by combining the ex-situ data collected in chapter six with ML and CV techniques. Specifically, the well-established Bag of Words approach was used to develop scale-invariant descriptions of in-situ melt pool morphology (Section 7.3). By combining unsupervised ML with fundamental knowledge of process space, links were made between unique in-situ melt pool morphologies and ex-situ outcomes such as *keyholing porosity* and *balling* (Section 7.4.1, Figures 7.22 – 7.31). Once determined, these linkages enabled the training of a melt pool morphology classification algorithm. Finally, the developed classification algorithm was used to study the exposure of non-bulk geometries including the edges of stripes (Section 7.4.2), unsupported overhangs (Section 7.4.3), and contours (Section 7.4.4).

8.2 Implications

Many of the results presented in this thesis have broader implications which are presented below. Implications relating to the expansion of process space are enumerated first, followed by implications relating to the development of in-situ process monitoring methodologies.

- 9. The collection of data from multiple AlSi10Mg melt pool cross-sections and the quantification of bulk porosity enabled the development of an AlSi10Mg L-PBF processing window that is more robust than those currently available in the literature.
- 10. Analysis of the statistical behavior of AlSi10Mg and In718 melt pool behavior indicates that while the geometry of the majority of melt pools follow a normal distribution, unexplained outliers do exist. Such behavior implies that selection of conservative processing parameters (e.g. laser beam power, travel velocity, and hatch spacing) based on simple variability metrics (e.g. standard deviation) may not be sufficient to ensure part quality (e.g. a lack of porosity). Indeed, this observed melt pool behavior further motivates the development of in-situ monitoring methodologies.
- 11. Comparison of In718 melt pools exposed on a powder layer to In718 melt pools exposed on a bare substrate [12] suggests that melt pool geometry, particularly cross-sectional depth, may be significantly altered in the presence of a powder layer. Similar behavior is also suggested by the comparison of measured AlSi10Mg lack-of-fusion porosity to that predicted by analytical models reported in the literature [72].
- 12. Study of multiple In718 melt pool cross-sections confirms, as has been reported in the literature, that certain processing flaws such as keyholing porosity and balling occur irregularly and therefore analysis of a single melt pool cross-section is often insufficient

to properly characterize melt pool morphology for a given process parameter combination. This work, combined with the bulk porosity analysis presented for AlSi10Mg, also suggests the existence of a stable keyholing regime at higher beam travel velocities. In this regime, keyhole-mode melting occurs without the formation of keyholing porosity. The existence of such a regime has been previously noted in the welding literature [84].

- 13. The successful manufacture of test artifacts using non-standard powders in an L-PBF process has the potential to dramatically improve the robustness of the feedstock supply chain and reduce part production costs [89]. Furthermore, while a correlation between mean powder particle size and worsening part quality was identified, overall part quality remained high for all of the tested Powder Systems. Bulk porosity did not exceed 0.1% for any of the powder systems and no trend was observed for one of the edge roughness metrics most likely to affect the fatigue life of a part (based on internal research at CMU by Christopher Kantzos).
- 14. The development of an L-PBF powder bed anomaly detection and classification algorithm paves the way for real-time feedback control of the process. Furthermore, the developed methodology does not rely on human-created heuristics and is therefore highly extensible to alternate material systems and other powder bed-based AM technologies.
- 15. The novel usage of a pre-trained CNN to analyze input data at multiple size scales in order to improve patch-wise classification accuracy is a promising new avenue for ML applications in the AM arena. Indeed, such an approach may also be effective in other

manufacturing applications and fields as disparate as microstructure analysis and medical imaging [151].

- 16. While initially developed for process monitoring applications, the presented powder bed anomaly classification algorithm has proven to be an extremely powerful data analytics tool. Used offline (i.e. post build-completion) it can be, and has been, used to understand complex build failure modes and inform the redesign of parts to improve the chance of success in subsequent builds. It has also demonstrated the potential to quickly correlate powder spreading anomalies and part deformation with fusion process parameters and laser scan strategy – enabling several new avenues of L-PBF research.
- 17. While initially developed strictly to enable the study of in-situ melt pool morphology with a fixed Field of View high speed camera, the presented coaxial transformation algorithm is broadly applicable. Specifically, it has the potential to benefit the work of other researchers studying L-PBF melt pool dynamics and has already been used by Fisher [34] to facilitate the measurement of melt pool temperature fields, cooling rates, and thermal gradients.
- 18. Unsupervised ML techniques applied to melt pool (emitted light) images captured using a high speed visible-light camera revealed the existence of multiple in-situ morphologies unique to certain regions of L-PBF process space. Furthermore, these unique morphologies can be tentatively linked to ex-situ flaws such as keyholing porosity and balling. Previously, work in the literature has only correlated in-situ melt pool morphologies with keyholing porosity in the LENS DED AM process using a non-scale invariant and spatter-agnostic description of in-situ melt pool morphology [202].

19. The successful linkage of in-situ and ex-situ melt pool morphologies enabled the usage of supervised ML techniques for melt pool classification. As preliminarily demonstrated in this thesis, this methodology allows for the study of fusion of non-bulk geometries and the identification of melt pool-scale flaws triggered by local build geometry. Such a capability would be invaluable for the future development of process parameters optimized for the production of defect-free non-bulk geometries.

8.3 Future Work

The results presented in this thesis suggest numerous exciting avenues for future research topics, many of which have the potential to significantly impact metal Additive Manufacturing for years to come. A selection of suggested future research topics is presented below.

- Analyze the geometry of AlSi10Mg melt pools exposed on a layer of powder in order to determine if the observed disagreement between the bulk porosity measurements and the literature lack-of-fusion model is due to a difference in melt pool size between the powder and no powder environments.
- 2. Sample additional melt pool cross-sections to better characterize the outlier behavior observed for the AlSi10Mg and In718 material systems. Such an analysis combined with in-situ monitoring of the melt pool may also allow for the identification of the root cause(s) of these outliers.
- Perform mechanical testing on artifacts built using non-standard powder systems in order to directly quantify the effects of powder particle size on as-built part performance.

- 4. Utilize a testbed L-PBF machine, such as the AMMT at the National Institutes of Standards and Technology (NIST) [223] to implement the presented powder bed anomaly detection and classification algorithm in a real-time environment. Once implemented, develop appropriate heuristics and material "re-work" strategies [133] in order to mitigate detected millimeter-scale flaws.
- Retrain the presented MsCNN algorithm to detect and classify powder bed anomalies in other powder bed-based AM processes such as Metal Binder Jetting and EB-PBF in order to further demonstrate the flexibility of the developed ML architecture.
- 6. The current CNN utilizes a three channel input layer one for each of the three size scales. A custom CNN architecture could be designed with many input channels in order to accept data not just at different size scales but also from different sensor modalities (e.g. a near-infrared camera) and different time periods (e.g. previous layers). The simultaneous analysis of these disparate data streams may improve anomaly detection and classification accuracies.
- 7. Utilize a higher resolution camera for monitoring the powder bed in order to enable detection of smaller-scale anomalies. Utilize multiple light sources to mitigate anomaly detection artifacts such as those observed in Section 3.3.2. Note that lighting should remain unidirectional for each powder bed image as many anomalies are visible primarily due to the casting of shadows; the results from multiple images captured under differing lighting conditions can then be combined.
- 8. Utilize the presented MsCNN to further investigate the observed correlation between *super-elevation* (and other powder bed anomalies) and fusion process parameters such

as hatch spacing and laser scan strategy. Utilize the presented MsCNN to study the behavior of difficult-to-spread powder systems such as the EOS standard AlSi10Mg powder and optimize process parameters such as the recoater blade travel speed in order to reduce the occurrence of powder bed anomalies.

- 9. The presented melt pool morphology classification algorithm relies on a combination of unsupervised ML and expert human input in order to train a classifier. As a result, while reasonably strong statements can be made regarding the statistical behavior of a large set of melt pool images, only weak statements can be made regarding the classification of an individual melt pool. Addressing this limitation requires the collection of training data connected to ground truth information in order to enable direct usage of supervised ML techniques. Collection of such a dataset may be possible by simultaneous above-view imaging of a melt pool using the presented high speed camera setup and side-view imaging of the melt pol using a Dynamic X-Ray (DXR) setup such as that at Argonne National Lab [85].
- 10. As discussed in Chapter 7, much of the *feature* extraction procedure used to describe insitu melt pool morphology remains un-optimized. If appropriate ground-truth information can be collected, Deep Learning techniques (such as those used in Chapter 4) may allow for the development of a substantially more robust description of in-situ melt pool morphology. Such methods could enable improved differentiation between in-situ melt pool appearances.
- 11. While in-situ melt pool morphology was studied for several non-bulk geometries in this thesis, an enormous number of non-bulk situations remain unexplored. Of particular

interest may be thin wall structures (exposed using a melt track traveling along the long dimension of the structure) and the influence of shielding gas flow with respect to melt pool travel direction. Furthermore, the presented works focuses primarily on the In718 EOS nominal parameter combination and does not explore process parameter optimization for non-bulk geometries.

References

- T. M. Rankin, N. A. Giovinco, D. J. Cucher, G. Watts, B. Hurwitz, and D. G. Armstrong, "Three-dimensional printing surgical instruments: Are we there yet?," *J. Surg. Res.*, vol. 189, no. 2, pp. 193–197, 2014.
- [2] D. Bourell, M. Leu, and D. Rosen, "Roadmap for additive manufacturing: Identifying the Future of Freeform Processing," in *Solid Freeform Fabrication Proceedings*, 2009.
- [3] E. Atzeni and A. Salmi, "Economics of additive manufacturing for end-usable metal parts," *Int. J. Adv. Manuf. Technol.*, vol. 62, no. 9–12, pp. 1147–1155, 2012.
- [4] D. S. Thomas and S. W. Gilbert, "Costs and Cost Effectiveness of Additive Manufacturing A Literature Review and Discussion," 2014.
- [5] K. M. Taminger and R. A. Hafley, "Electron beam freeform fabrication for cost effective near-net shape manufacturing," *NATO/RTOAVT-139 Spec. Metting Cost Eff. Manuf. via Net Shape Process.*, 2006.
- [6] T. Wohlers and T. Caffrey, "Wohlers Report," Fort Collins, CO, 2011.
- [7] C. A. Giffi, B. Gangula, and P. Illinda, "3D opportunity for the automotive industry: Additive manufacturing hits the road," 2014.
- [8] C. Emmelmann, P. Sander, J. Kranz, and E. Wycisk, "Laser Additive Manufacturing and Bionics: Redefining Lightweight Design," *Phys. Procedia*, vol. 12, pp. 364–368, 2011.
- [9] D. A. Ramirez *et al.*, "Open-cellular copper structures fabricated by additive manufacturing using electron beam melting," *Mater. Sci. Eng. A*, vol. 528, no. 16–17, pp. 5379–5386, Jun. 2011.
- [10] T. Kellner, "The FAA Cleared the First 3D Printed Part to Fly in a Commercial Jet Engine from GE," *GE Reports*, Aug-2015.
- [11] D. C. Hofmann *et al.*, "Developing gradient metal alloys through radial deposition additive manufacturing.," *Sci. Rep.*, 2014.
- [12] S. P. Narra, "Melt Pool Geometry and Microstructure Control Across Alloys in Metal Based Additive Manufacturing Processes," Carnegie Mellon University, 2017.
- [13] S. H. Khajavi, J. Partanen, and J. Holmströ, "Additive manufacturing in the spare parts supply chain," *Comput. Ind.*, vol. 65, pp. 50–63, 2014.
- [14] M. Gravier, "3D Printing: Customers taking charge of the supply chain," *IndustryWeek*, Apr-2016.
- [15] S. H. Huang, P. Liu, A. Mokasdar, and L. Hou, "Additive manufacturing and its societal impact: a literature review," Int. J. Adv. Manuf. Technol., vol. 67, no. 5–8, pp. 1191–1203, Jul. 2013.
- U. S. Navy, "Shipboard additive manufacturing (AM)/3D printing SBIR 2016.1 N161-038," 2016. [Online]. Available: http://www.navysbir.com/n16_1/N161-038.htm. [Accessed: 01-Jan-2016].

- [17] 3D Hubs, "Additive Manufacturing Technologies," 2016. [Online]. Available: https://www.3dhubs.com/what-is-3d-printing#additive-manufacturing-infographic.
- [18] ExOne, "Technology Overview." [Online]. Available: http://www.exone.com/Resources/Technology-Overview. [Accessed: 04-Jan-2018].
- [19] Desktop Metal, "Office-friendly metal 3D printing." [Online]. Available: https://www.desktopmetal.com/products/studio/. [Accessed: 04-Jan-2018].
- [20] M. Kerschbaumer and G. Ernst, "Hybrid manufacturing process for rapid high performance tooling combining high speed milling and laser cladding," in *Proceedings of the 23rd International Congress on Applications of Lasers and Electro-Optics (ICALEO), San Francisco, CA*, 2004, vol. 97, pp. 1710–1720.
- [21] Sciaky, "Electron Beam Manufacturing Technology," 2018. [Online]. Available: http://www.sciaky.com/additive-manufacturing/electron-beam-additive-manufacturingtechnology. [Accessed: 04-Jan-2018].
- [22] M. Kottman, S. Zhang, J. McGuffin-Cawley, P. Denney, and B. K. Narayanan, "Laser Hot Wire Process: A Novel Process for Near-Net Shape Fabrication for High-Throughput Applications," J. Mater., vol. 67, no. 3, pp. 622–628, 2015.
- [23] TWI, "Welding Automation," 2018. [Online]. Available: https://www.twiglobal.com/capabilities/joining-technologies/welding-engineering/welding-automation/. [Accessed: 04-Jan-2018].
- [24] Arcam, "Electron Beam Melting." [Online]. Available: http://www.arcam.com/wpcontent/uploads/arcamebm-corp-brochure-fnlv3.pdf. [Accessed: 04-Jan-2018].
- [25] EOS GmbH, "Additive Manufacturing, Laser-Sintering and Industrial 3D Printing Benefits and Functional Principle." [Online]. Available: https://www.eos.info/additive_manufacturing/for_technology_interested?_scrivito_disp lay_mode=view&_scrivito_workspace_id=published. [Accessed: 04-Jan-2018].
- [26] I. Gibson, D. W. Rosen, and B. Stucker, *Additive Manufacturing Technologies: Rapid Prototyping to Direct Digital Manufacturing*. Springer, 2010.
- [27] EOS GmbH, "Materials for Metal Additive Manufacturing." [Online]. Available: https://www.eos.info/material-m.
- [28] Arcam, "Metal Powders." [Online]. Available: http://www.arcam.com/technology/products/metal-powders/. [Accessed: 04-Jan-2018].
- [29] J. Beuth *et al.*, "Process Mapping for Qualification Across Multiple Direct Metal Additive Manufacturing Processes," in *Solid Freeform Fabrication Proceedings*, 2013, pp. 655–665.
- [30] J. Fox, "Transient Melt Pool Response in Additive Manufacturing of Ti-6Al-4V," Carnegie Mellon University, 2015.
- [31] Z. Francis, "The Effects of Laser and Electron Beam Spot Size in Additive Manufacturing Processes," Carnegie Mellon University, 2017.

- [32] J. Beuth and N. Klingbeil, "The role of process variables in laser-based direct metal solid freeform fabrication," J. Mater., vol. 53, no. 9, pp. 36–39, 2001.
- [33] D. Clymer, J. Beuth, and J. Cagan, "Additive Manufacturing Process Design," in *Solid Freeform Fabrication Proceedings*, 2016.
- [34] B. Fisher, "Part Temperature Effects in Powder Bed Fusion Additive Manufacturing of Ti-6Al-4V," Carnegie Mellon University, 2018.
- [35] B. Fisher, J. Mireles, S. Ridwan, R. Wicker, and J. Beuth, "Consequences of Part Temperature Variability in Electron Beam Melting of Ti-6Al-4V," *J. Mater.*, 2017.
- [36] C. Montgomery, "The Effect of Alloys, Powder, and Overhanging Geometries in Laser Powder Bed Additive Manufacturong," Carnegie Mellon University, 2017.
- [37] J. Fox and J. Beuth, "Process Mapping of Transient Melt Pool Response in Wire Feed E-Beam Additive Manufacturing of Ti-6Al-4V," Proc. Solid Free. Fabr. Symp., pp. 675–683, 2013.
- [38] M. Lansard, "The complete list of metal additive manufacturing companies and metal 3D printer brands," *Aniwaa*, Sep-2016.
- [39] M. Matthews, G. Guss, S. Khairallah, A. Rubenchick, P. Depond, and W. King, "Denudation of metal powder layers in laser bed fusion," *Acta Mater.*, vol. 114, pp. 33– 42, 2016.
- [40] EOS GmbH, "EOS M 290 Data Sheet," 2017. [Online]. Available: https://cdn0.scrvt.com/eos/413c861f2843b377/51d966b8b31d/EOS_System_Data_Shee t_EOS_M_290_EN_V6_Web.pdf. [Accessed: 20-Jan-2018].
- [41] A. Anam, J. J. S. Dilip, D. Pal, and B. Stucker, "Effect of scan pattern on the microstructural evolution of Inconel 625 during Selective Laser Melting," in *Solid Freeform Fabrication*, 2014, pp. 363–376.
- [42] E. Soylemez, J. L. Beuth, and K. Taminger, "Controlling Melt Pool Dimensions over a Wide Range of Material Deposition Rates in Electron Beam Additive Manufacturing," *Solid Free. Fabr.*, pp. 571–582, 2010.
- [43] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [44] G. Csurka, C. Dance, L. Fan, J. Willamowski, and B. Cedric, "Visual categorization with bags of keypoints," in Workshop on Statistical Learning in Computer Vision (ECCV), 2004, pp. 1–22.
- [45] EOS GmbH, "EOS Software for Additive Manufacturing." [Online]. Available: https://www.eos.info/systems_solutions/software. [Accessed: 31-May-2017].
- [46] O. Hedin, "Production Technology for Additive Manufacturing The Arcam EBM [®] process: A walkthrough," in *OMTEC Proceedings*, 2012.

- [47] M. Seifi, D. Christiansen, J. Beuth, O. Harrysson, and J. J. Lewandowski, "Process Mapping, Fracture and Fatigue Behaviour of Ti-6Al-4V Produced by EBM Additive Manufacturing," in *Ti-2015-The 13th World Conference on Titanium*, 2016, pp. 1373– 1378.
- [48] R. Cunningham, S. P. Narra, C. Montgomery, J. Beuth, and A. D. Rollett, "Synchrotron-Based X-Ray Microtomography Characterization of the Effect of Processing Variables on Porosity Formation in Laser Power-Bed Additive Manufacturing of Ti-6Al-4V," J. Mater., vol. 69, 2016.
- [49] P. A. Kobryn and S. L. Semiatin, "Microstructure and texture evolution during solidification processing of Ti–6Al–4V," J. Mater. Process. Technol., vol. 135, no. 2–3, pp. 330–339, 2003.
- [50] S. Tammas-Williams and I. Todd, "Design for additive manufacturing with site-specific properties in metals and alloys," *Scr. Mater.*, vol. 135, pp. 105–110, 2016.
- [51] A. Leandri, "A Look into Powder Materials for Metal 3D Printing," *3D Printing Industry*, Sep-2015.
- [52] R. E. Laureijs, J. Bonnín Roca, S. Prabha Narra, C. Montgomery, J. L. Beuth, and E. R. H. Fuchs, "Metal Additive Manufacturing: Cost Competitive Beyond Low Volumes," J. Manuf. Sci. Eng., vol. 139, no. 8, pp. 1–9, 2016.
- [53] J. Beuth *et al.*, "A Database Relating Powder Properties to Process Outcomes for Direct Metal AM," 2017.
- [54] L. Scime and J. Beuth, "Anomaly Detection and Classification in a Laser Powder Bed Additive Manufacturing Process using a Trained Computer Vision Algorithm," Addit. Manuf., vol. 19, pp. 114–126, 2018.
- [55] L. Scime, B. Fisher, and J. Beuth, "Using coordinate transforms to improve the utility of a fixed field of view high speed camera for additive manufacturing applications," *Manuf. Lett.*, vol. 15PB, pp. 104–106, 2018.
- [56] ASTM, "Standard Terminology for Additive Manufacturing." 2015.
- [57] EOS GmbH, "Material data sheet for EOS Aluminium AlSi10Mg." 2014.
- [58] K. Kempen, L. Thijs, E. Yasa, M. Badrossamay, W. Verheecke, and J. P. Kruth, "Process Optimization and Micostructural Analysis for Selective Laser Melting of AlSi10Mg," *Solid Free. Fabr.*, pp. 484–495, 2011.
- [59] Lytron, "Defense & Aerospace Thermal Solutions," 2017. [Online]. Available: http://www.lytron.com/Industries/Defense-And-Aerospace-Thermal-Solutions. [Accessed: 14-Jun-2017].
- [60] AAVID Thermacore, "Thermal Management for Aircraft, Spacecraft and Satellites." [Online]. Available: http://www.thermacore.com/industries/aerospace-avionics.aspx. [Accessed: 14-Jun-2017].
- [61] M. E. Ellion, J. W. Putt, and J. J. Montich, "Liquid-vapor separator," US 4272257 A, 1976.

- [62] B. A. Fulcher, D. K. Leigh, and T. J. Watt, "Comparison of AlSi10Mg and Al 6061 Processed Through DMLS," in *Solid Freeform Fabrication*, 2014, pp. 404–419.
- [63] T. Mukherjee, J. S. Zuback, A. De, and T. Debroy, "Printability of alloys for additive manufacturing," *Nat. Sci. Reports*, 2015.
- [64] Aluminum Brazing, "Cladding Alloys," *News, Knowledge & Technology*, 2010. [Online]. Available: http://www.aluminium-brazing.com/2010/09/. [Accessed: 12-Jun-2017].
- [65] E. Brandl, U. Heckenberger, V. Holzinger, and D. Buchbinder, "Additive manufactured AlSi10Mg samples using Selective Laser Melting (SLM): Microstructure, high cycle fatigue, and fracture behavior," *Mater. Des.*, vol. 34, pp. 159–169, 2012.
- [66] A. Vasinonta and J. Beuth, "Process Maps for Controlling Residual Stress and Melt Pool Size in Laser-Based SFF Processes," in *Solid Freeform Fabrication*, 2000, pp. 200–208.
- [67] P. Promoppatum, R. Onler, and S. C. Yao, "Numerical and experimental investigations of micro and macro characteristics of direct metal laser sintered Ti-6Al-4V products," J. Mater. Process. Technol., vol. 240, pp. 262–273, 2017.
- [68] C. Montgomery, J. Beuth, L. Sheridan, and N. Klingbeil, "Process Mapping of Inconel 625 in Laser Powder Bed Additive Manufacturing," in *Solid Freeform Fabrication Proceedings*, 2015.
- [69] N. T. Aboulkhair, I. Maskery, C. Tuck, I. Ashcroft, and N. M. Everitt, "On the formation of AlSi10Mg single tracks and layers in selective laser melting: Microstructure and nanomechanical properties," J. Mater. Process. Technol., vol. 230, pp. 88–98, 2016.
- [70] N. Read, W. Wang, K. Essa, and M. M. Attallah, "Selective laser melting of AlSi10Mg alloy: Process optimisation and mechanical properties development," *Mater. Des.*, vol. 65, pp. 417–424, 2015.
- [71] M. Tang and P. C. Pistorius, "Oxides, porosity and fatigue performance of AlSi10Mg parts produced by selective laser melting," *Int. J. Fatigue*, vol. 94, pp. 192–201, 2017.
- [72] M. Tang, P. C. Pistorius, and J. L. Beuth, "Prediction of lack-of-fusion porosity for powder bed fusion," *Addit. Manuf.*, vol. 14, pp. 39–48, 2017.
- [73] P. S. Wei, K. C. Chuang, J. S. Ku, and T. Debroy, "Mechanisms of spiking and humping in keyhole welding," *IEEE Trans. Components, Packag. Manuf. Technol.*, vol. 2, no. 3, pp. 383–394, 2012.
- [74] H. Gong, K. Rafi, H. Gu, T. Starr, and B. Stucker, "Analysis of defect generation in Ti–6Al–
 4V parts made using powder bed fusion additive manufacturing processes," Addit. Manuf., vol. 1–4, no. 4, pp. 87–98, 2014.
- [75] EOS GmbH, "EOS M 280 Data Sheet," 2013. [Online]. Available: https://cdn0.scrvt.com/eos/e1dc925774b24d9f/8d4013a5bc26/EOS_System_Data_Shee t_EOSINT_M280_EN_V2_Web.pdf. [Accessed: 20-Jan-2018].
- [76] Makino, "Advanced Techniques in Aerospace Manufacturing From Makino," 2015. [Online]. Available: http://www.makino.com/industries/edm/. [Accessed: 09-Feb-2017].

- [77] ASTM, "E3-11 Standard Guide for Preparation of Metallographic Specimens 1," ASTM. 2011.
- [78] US National Institutes of Health, "ImageJ." [Online]. Available: https://imagej.nih.gov/ij/. [Accessed: 13-Jun-2017].
- [79] R. Cunningham, S. P. Narra, T. Ozturk, J. Beuth, and A. D. Rollett, "Evaluating the Effect of Processing Parameters on Porosity in Electron Beam Melted Ti-6Al-4V via Synchrotron Xray Microtomography," J. Mater., vol. 68, no. 3, pp. 765–771, 2016.
- [80] A. McDonald, M. Lamontagne, S. Chandra, and C. Moreau, "Photographing Impact of Plasma-Sprayed Particles on Metal Substrates," in *International Thermal Spray Conference*, 2006.
- [81] S. Clijsters, T. Craeghs, S. Buls, K. Kempen, and J. P. Kruth, "In situ quality control of the selective laser melting process using a high-speed, real-time melt pool monitoring system," Int. J. Adv. Manuf. Technol., vol. 75, pp. 1089–1101, 2014.
- [82] M. Baron, Probability and Statistics for Computer Scientists, Second. Boca Raton, FL: CRC Press, 2016.
- [83] J. W. Elmer, W. H. Giedt, and T. W. Eager, "The Transition from Shallow to Deep Penetration during Electron Beam Welding," *Weld. J.*, vol. 69, no. 5, pp. 167–176, 1990.
- [84] D. T. Pastor M, Zhao H, Martukanitz RP, "Porosity, Underfill and Magnesium Loss during Continuous Wave Nd : YAG Laser Welding of Thin Plates of Aluminum Alloys 5182 and 5754," Weld. J. - Incl. Weld. Res. Suppl., vol. 78, p. 207s–216s, 1999.
- [85] C. Zhao *et al.*, "Real-time monitoring of laser powder bed fusion process using highspeed X-ray imaging and diffraction," *Sci. Rep.*, 2017.
- [86] P. Edwards, A. O'Conner, and M. Ramulu, "Electron Beam Additive Manufacturing of Titanium Components: Properties and Performance," J. Manuf. Sci. Eng., vol. 135, no. 6, 2013.
- [87] U. Gratzke, P. D. Kapadia, J. Dowden, J. Kroos, and G. Simon, "Theoretical approach to the humping phenomenon in welding processes," J. Phys. D. Appl. Phys., vol. 25, no. 11, pp. 1640–1647, 2000.
- [88] I. Yadroitsev, A. Gusarov, I. Yadroitsava, and I. Smurov, "Single track formation in selective laser melting of metal powders," J. Mater. Process. Technol., vol. 210, no. 12, pp. 1624–1631, Sep. 2010.
- [89] U. Fritsching and V. Uhlenwinkel, *Powder Metallurgy*. InTech, 2012.
- [90] J. Dawes, R. Bowerman, and R. Trepleton, "Introduction to the Additive Manufacturing Powder Metallurgy Supply Chain," *Johnson Matthey Technol. Rev.*, vol. 59, no. 3, pp. 243–256, 2015.
- [91] S. Froes, "Titanium Powder Metallurgy: A Review Part 1," Adv. Mater. Process., 2012.

- [92] G. Jacob, A. Donmez, J. Slotwinski, and S. Moylan, "Measurement of powder bed density in powder bed fusion additive manufacturing processes," *Meas. Sci. Technol.*, vol. 27, 2016.
- [93] E. B. Herbold, O. Walton, and M. A. Homel, "Simulation of Powder Layer Deposition in Additive Manufacturing Processes Using the Discrete Element Method," 2015.
- [94] E. J. R. Parteli, "DEM simulation of particles of complex shapes using the multisphere method: Application for additive manufacturing," in *AIP Conference Proceedings*, 2013, pp. 185–188.
- [95] P. Desai, "Tribosurface Interactions involving Particulate Media with DEM-Calibrated Properties: Experiments and Modeling," Carnegie Mellon University, 2017.
- [96] A. B. Spierings, N. Herres, and G. Levy, "Influence of the particle size distribution on surface quality and mechanical properties in additive manufactured stainless steel parts," *Solid Free. Fabr.*, pp. 397–406, 2010.
- [97] D. B. Saint John, S. B. Joshi, T. W. Simpson, M. Qu, J. D. Rowatt, and Y. Lou, "A Preliminary Examination of Variability Due to Build Location and Powder Feedstock in Additive Manufacture of Inconel 718 using Laser-Based Powder Bed Fusion," in *Solid Freeform Fabrication*, 2016.
- [98] H. Gu et al., "Effects of Powder Variation on the Microstructure and Tensile Strength of Ti6Al4V Parts Fabricated by Selective Laser Melting," in Solid Freeform Fabrication Symposium, 2014, pp. 470–483.
- [99] National Institute of Standards and Technology, "NIST Additive Manufacturing Test Artifact," 2013. [Online]. Available: https://www.nist.gov/el/intelligent-systems-division-73500/production-systems-group/nist-additive-manufacturing-test. [Accessed: 27-May-2017].
- [100] W. E. King *et al.*, "Laser powder bed fusion additive manufacturing of metals; physics, computational, and materials challenges," *Appl. Phys. Rev.*, vol. 2, no. 4, p. 41304, 2015.
- [101] EOS GmbH, "Machine and software parameters EOSINT M 280." Munich, Germany, 2013.
- [102] EOS GmbH, "Material data sheet for EOS NickelAlloy IN718," 2014. [Online]. Available: http://ip-saas-eoscms.s3.amazonaws.com/public/4528b4a1bf688496/ff974161c2057e6df56db5b67f0f559 5/EOS_NickelAlloy_IN718_en.pdf.
- [103] S. A. Saltykov, Stereometric Metallography. Moscow, 1958.
- [104] J. Gegner, "2D-3D Conversion of Object Size Distributions in Quantitative Metallography," in *Mathematical Modeling and Computer Simulation of Metal Technologies*, 2006.
- [105] EOS GmbH, "EOS monitoring solutions available for EOS M290." [Online]. Available: https://www.eos.info/eos-m290. [Accessed: 31-May-2017].

- [106] A. Ladewig, G. Schlick, M. Fisser, V. Schulze, and U. Glatzel, "Influence of the shielding gas flow on the removal of process by-products in the selective laser melting process," *Addit. Manuf.*, vol. 10, pp. 1–9, 2016.
- [107] S. A. Khairallah, A. T. Anderson, A. Rubenchik, and W. E. King, "Laser powder-bed fusion additive manufacturing: Physics of complex melt flow and formation mechanisms of pores, spatter, and denudation zones," *Acta Mater.*, vol. 108, pp. 36–45, 2016.
- [108] University of the West of England, "Pearson's Correlation Coefficient," 2018. [Online]. Available: http://learntech.uwe.ac.uk/da/Default.aspx?pageid=1442. [Accessed: 13-Jan-2018].
- [109] M. Grasso and B. M. Colosimo, "Process defects and *in situ* monitoring methods in metal powder bed fusion: a review," *Meas. Sci. Technol.*, vol. 28, no. 4, p. 44005, 2017.
- [110] S. K. Everton, M. Hirsch, P. Stravroulakis, R. K. Leach, and A. T. Clare, "Review of in-situ process monitoring and in-situ metrology for metal additive manufacturing," *JMADE*, vol. 95, pp. 431–445, 2016.
- [111] T. G. Spears and S. A. Gold, "In-process sensing in selective laser melting (SLM) additive manufacturing," *Integr. Mater. Manuf. Innov.*, vol. 5, no. 2, pp. 1–25, 2016.
- [112] T. Furumoto, T. Ueda, M. R. Alkahari, and A. Hosokawa, "Investigation of laser consolidation process for metal powder by two-color pyrometer and high-speed video camera," *CIRP Ann. - Manuf. Technol.*, vol. 62, no. 1, pp. 223–226, 2013.
- [113] M. Islam, T. Purtonen, H. Piili, A. Salminen, and O. Nyrhilä, "Temperature profile and imaging analysis of laser additive manufacturing of stainless steel," *Phys. Procedia*, vol. 41, pp. 835–842, 2013.
- [114] B. Lane, S. Moylan, E. Whitenton, and L. Ma, "Thermographic Measurements of the Commercial Laser Powder Bed Fusion Process at NIST," *Proc. Solid Free. Fabr. Symp.*, pp. 575–591, 2015.
- [115] J. Raplee *et al.*, "Thermographic Microstructure Monitoring in Electron Beam Additive Manufacturing," *Sci. Rep.*, vol. 7, pp. 1–16, 2017.
- [116] H. Krauss, T. Zeugner, and M. F. Zaeh, "Layerwise monitoring of the Selective Laser Melting process by thermography," *Phys. Procedia*, vol. 56, no. C, pp. 64–71, 2014.
- [117] S. Yoder *et al.,* "Characterization of Topology Optimized Ti-6Al-4V Components using Electron Beam Powder Bed Fusion," *Addit. Manuf.*, vol. 19, pp. 184–196, 2017.
- [118] A. J. Dunbar *et al.*, "Development of experimental method for in situ distortion and temperature measurements during the laser powder bed fusion additive manufacturing process," *Addit. Manuf.*, vol. 12, pp. 25–30, 2016.
- [119] A. R. Nassar et al., "Sensing for directed energy deposition and powder bed fusion additive manufacturing at Penn State University," Laser 3D Manuf. III SPIE, vol. 9738, 2016.

- [120] B. Lane, S. Mekhontsev, S. Grantham, M. Vlasea, J. Whiting, and H. Yeun, "Design, Developments, and Results from the NIST Additive Manufacturing Metrology Testbed (AMMT)," in Solid Freeform Fabrication, 2016, pp. 1145–1160.
- [121] G. Repossini, V. Laguzza, M. Grasso, and B. M. Colosimo, "On the use of spatter signature for in-situ monitoring of Laser Powder Bed Fusion," *Addit. Manuf.*, vol. 16, pp. 35–48, 2017.
- [122] M. Grasso, A. G. Demir, B. Previtali, and B. M. Colosimo, "In situ monitoring of selective laser melting of zinc powder via infrared imaging of the process plume," *Robot. Comput. Integr. Manuf.*, vol. 49, pp. 229–239, 2018.
- [123] J. C. Heigel and B. M. Lane, "Measurement of the Melt Pool Length During Single Scan Tracks in a Commercial Laser Powder Bed Fusion Process," in ASME International Manufacturing Science and Engineering Conference, 2017.
- [124] M. Grasso, V. Laguzza, Q. Semeraro, and B. M. Colosimo, "In-Process Monitoring of Selective Laser Melting: Spatial Detection of Defects Via Image Data Analysis," J. Manuf. Sci. Eng., vol. 139, no. 5, 2016.
- [125] P. Bidare, R. R. J. Maier, R. J. Beck, J. D. Shephard, and A. J. Moore, "An open-architecture metal powder bed fusion system for in-situ process measurements," *Addit. Manuf.*, vol. 16, pp. 177–185, Aug. 2017.
- [126] J. A. Kanko, A. P. Sibley, and J. M. Fraser, "In situ morphology-based defect detection of selective laser melting through inline coherent imaging," J. Mater. Process. Technol., vol. 231, pp. 488–500, 2016.
- [127] E. R. Denlinger, J. C. Heigel, P. Michaleris, and T. A. Palmer, "Effect of inter-layer dwell time on distortion and residual stress in additive manufacturing of titanium and nickel alloys," J. Mater. Process. Technol., vol. 215, pp. 123–131, 2015.
- [128] H. Rieder, A. Dillhöfer, M. Spies, J. Bamberg, and T. Hess, "Online Monitoring of Additive Manufacturing Processes Using Ultrasound," in *Proceedings of the 11th European Conference on Non-Destructive Testing*, 2014, vol. 11, pp. 2194–2201.
- [129] J. A. Slotwinski and E. J. Garboczi, "Porosity of additive manufacturing parts for process monitoring," *AIP Conf. Proc.*, vol. 1581 33, no. 2014, pp. 1197–1204, 2014.
- [130] S. A. Shevchik, C. Kenel, C. Leinenbach, and K. Wasmer, "Acoustic emission for in situ quality monitoring in additive manufacturing using spectral convolutional neural networks," *Addit. Manuf.*, 2017.
- [131] S. Kleszczynski, J. zur Jacobsmühlen, B. Reinarz, J. T. Sehrt, G. Witt, and D. Merhof, "Improving Process Stability of Laser Beam Melting Systems," in *Fraunhofer Direct Digital Manufacturing Conference (DDMC)*, 2014, pp. 1–6.
- [132] R. J. Smith, M. Hirsch, R. Patel, W. Li, A. T. Clare, and S. D. Sharples, "Spatially resolved acoustic spectroscopy for selective laser melting," *J. Mater. Process. Technol.*, vol. 236, pp. 93–102, 2016.

- [133] M. Hirsch *et al.*, "Targeted rework strategies for powder bed additive manufacture," *Addit. Manuf.*, vol. 19, pp. 127–133, 2018.
- [134] J. Zur Jacobsmuhlen, J. Achterhold, S. Kleszczynski, G. Witt, and D. Merhof, "Robust calibration marker detection in powder bed images from laser beam melting processes," in *Proceedings of the IEEE International Conference on Industrial Technology*, 2016, pp. 910–915.
- [135] S. Kleszczynski, J. Zur Jacobsmühlen, J. T. Sehrt, and G. Witt, "Error detection in laser beam melting systems by high resolution imaging," *Solid Free. Fabr.*, pp. 975–987, 2012.
- [136] T. Craeghs, S. Clijsters, E. Yasa, and J.-P. Kruth, "Online quality control of selective laser melting," in *Solid Freeform Fabrication Proceedings*, 2011, pp. 212–226.
- [137] J. zur Jacobsmuhlen, S. Kleszczynski, G. Witt, and D. Merhof, "Robustness analysis of imaging system for inspection of laser beam melting systems.," 2014 IEEE Emerg. Technol. Fact. Autom., pp. 1–4, 2014.
- [138] M. Aminzadeh, "A Machine Vision System for In-Situ Quality Inspection in Metal Powder-Bed Additive Manufacturing," Georgia Institute of Technology, 2016.
- [139] J. zur Jacobsmuhlen, S. Kleszczynski, D. Schneider, and G. Witt, "High resolution imaging for inspection of laser beam melting systems," in *IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, 2013, pp. 707–712.
- [140] J. Zur Jacobsmuhlen, S. Kleszczynski, G. Witt, and D. Merhof, "Detection of elevated regions in surface images from laser beam melting processes," in *Conference of the IEEE Industrial Electronics Society (IECON)*, 2016, pp. 1270–1275.
- [141] B. K. Foster, E. W. Reutzel, A. R. Nassar, B. T. Hall, S. W. Brown, and C. J. Dickman, "Optical, layerwise monitoring of powder bed fusion," in *Solid Freeform Fabrication*, 2015, pp. 295–307.
- [142] M. Abdelrahman, E. W. Reutzel, A. R. Nassar, and T. L. Starr, "Flaw detection in powder bed fusion using optical imaging," Addit. Manuf., vol. 15, pp. 1–11, 2017.
- [143] B. Zhang, J. Ziegert, F. Farahi, and A. Davies, "In situ surface topography of laser powder bed fusion using fringe projection," *Addit. Manuf.*, vol. 12, pp. 100–107, 2016.
- [144] A. L. Cooke and S. P. Moylan, "Process intermittent measurement for powder-bed based additive manufacturing," *Solid Free. Fabr.*, pp. 81–98, 2011.
- [145] C. Scott, "EOS Introduces EOSTATE Exposure OT, First Commercial Optical Tomography System for Additive Manufacturing," *3DPrint.com*, Jun-2017.
- [146] ConceptLaser, "Quality Management," 2018. [Online]. Available: https://www.conceptlaser.de/en/products/quality-management.html. [Accessed: 08-Feb-2018].
- [147] W. J. Sames, "Additive manufacturing of inconel 718 using electron beam melting: processing, post-processing, & mechanical properties," Texas A&M University, 2015.

- [148] J. Xiao, J. Hays, K. A. Ehinger, and A. Torralba, "SUN database: Large-scale scene recognition from abbey to zoo," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 3485–3492.
- [149] J. Winn, A. Criminisi, and T. Minka, "Object categorization by learned univeral dictionary," in *International Conference on Computer Vision (ICCV)*, 2005.
- [150] H. G. R. Gouk and A. M. Blake, "Fast Sliding Window Classification with Convolutional Neural Networks," Proc. 29th Int. Conf. Image Vis. Comput. New Zeal. - IVCNZ '14, pp. 114–118, 2014.
- [151] W. Shen, M. Zhou, F. Y. B, C. Yang, and J. T. B, "Multi-scale Convolutional Neural Networks for Lung Nodule Classification," in *International Conference on Information Processing in Medical Imaging*, 2015, vol. 9123, pp. 588–599.
- [152] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?," in *Neural Information Processing Systems Conference*, 2014, pp. 1–9.
- [153] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Adv. Neural Inf. Process. Syst.*, pp. 1–9, 2012.
- [154] E. Dubrofsky, "Homography estimation," The University of Vancouver, 2007.
- [155] ISO International, "ISO-12233:2014 Photography -- Electronic still picture imaging --Resolution and spatial frequency responses." Switzerland, 2014.
- [156] MathWorks, "imadjust," 2018. [Online]. Available: https://www.mathworks.com/help/images/ref/imadjust.html. [Accessed: 20-Oct-2018].
- [157] K. Zeng, N. Patil, H. Gu, and H. Gong, "Layer by layer validation of geometrical accuracy in additive manufacturing processes," in *Solid Freeform Fabrication*, 2013, pp. 76–87.
- [158] Autolt Consulting Ltd., "Autolt," 2016. [Online]. Available: https://www.autoitscript.com/site/autoit/. [Accessed: 20-Jan-2018].
- [159] MathWorks, "imresize," 2018. [Online]. Available: https://www.mathworks.com/help/images/ref/imresize.html. [Accessed: 20-Jan-2018].
- [160] R. Vaidya and S. Anand, "Image processing assisted tools for pre- and post-processing operations in additive manufacturing," *Procedia Manuf.*, vol. 5, pp. 958–973, 2016.
- [161] MathWorks, "imfilter," 2017. [Online]. Available: https://www.mathworks.com/help/images/ref/imfilter.html?s_tid=gn_loc_drop. [Accessed: 04-Jan-2017].
- [162] MathWorks, "gabor," 2017. [Online]. Available: https://www.mathworks.com/help/images/ref/gabor.html. [Accessed: 04-Jul-2017].
- [163] MathWorks, "k-Means Clustering," 2016. [Online]. Available: https://www.mathworks.com/help/stats/k-means-clustering.html. [Accessed: 30-Dec-2016].

- [164] MathWorks, "bsxfun," 2017. [Online]. Available: https://www.mathworks.com/help/matlab/ref/bsxfun.html. [Accessed: 04-Jul-2017].
- [165] M. K. Smith, "Overfitting," The University of Texas at Austin. [Online]. Available: http://www.ma.utexas.edu/users/mks/statmistakes/ovefitting.html. [Accessed: 30-Dec-2016].
- [166] MathWorks, "alexnet," 2018. [Online]. Available: https://www.mathworks.com/help/nnet/ref/alexnet.html. [Accessed: 05-Feb-2018].
- [167] MathWorks, "padarray," 2018. [Online]. Available: https://www.mathworks.com/help/images/ref/padarray.html. [Accessed: 05-Feb-2018].
- [168] A. Karpathy, "CS231n Convolutional Neural Networks for Visual Recognition," Standford Computer Science, 2017. [Online]. Available: http://cs231n.github.io/convolutionalnetworks/#overview. [Accessed: 05-Feb-2018].
- [169] Stanford Vision Lab, "ImageNet," 2016. [Online]. Available: http://www.image-net.org/. [Accessed: 05-Feb-2018].
- [170] W. Knight, "The Dark Secret at the Heart of AI," *MIT Technology Review*, vol. 120, no. 3, 2017.
- [171] R. Hahnloser, R. Sarpeshkar, M. Mahowald, R. Douglas, and S. Seung, "Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit," *Nature*, vol. 405, pp. 947–951, 2000.
- [172] P. Joshi, "What is Local Response Normalization in Convolutional Neural Networks," *Perpetual Enigma*, 2016. [Online]. Available: https://prateekvjoshi.com/2016/04/05/what-is-local-response-normalization-inconvolutional-neural-networks/. [Accessed: 06-Feb-2018].
- [173] MathWorks, "DropoutLayer," 2018. [Online]. Available: https://www.mathworks.com/help/nnet/ref/nnet.cnn.layer.dropoutlayer.html. [Accessed: 06-Feb-2018].
- [174] MathWorks, "trainSoftmaxLayer," 2018. [Online]. Available: https://www.mathworks.com/help/nnet/ref/trainsoftmaxlayer.html. [Accessed: 07-Feb-2018].
- [175] A. Deshpande, "A Beginner's Guide to Understanding Convolutional Neural Networks," 2016. [Online]. Available: https://adeshpande3.github.io/adeshpande3.github.io/A-Beginner's-Guide-To-Understanding-Convolutional-Neural-Networks/. [Accessed: 07-Feb-2018].
- [176] A. Ng et al., "Optomization: Stochastic Gradient Descent." [Online]. Available: http://ufldl.stanford.edu/tutorial/supervised/OptimizationStochasticGradientDescent/. [Accessed: 07-Feb-2018].

- [177] MathWorks, "trainingOptions," 2018. [Online]. Available: https://www.mathworks.com/help/nnet/ref/trainingoptions.html. [Accessed: 07-Feb-2018].
- [178] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, pp. 861– 874, 2005.
- [179] J. Brownlee, "What is the Difference Between Test and Validation Datasets," Machine Learning Mastery, 2017. [Online]. Available: https://machinelearningmastery.com/difference-test-validation-datasets/. [Accessed: 08-Feb-2018].
- [180] M. Kuhn and K. Johnson, *Applied Predictive Modeling*, 5th ed. New York: Springer US, 2016.
- [181] MathWorks, "MATLAB Runtime," 2017. [Online]. Available: https://www.mathworks.com/products/compiler/mcr.html. [Accessed: 10-Jun-2017].
- [182] J. Zur Jacobsmühlen, S. Kleszczynski, G. Witt, and D. Merhof, "Elevated Region Area Measurement For Quantitative Analysis Of Laser Beam Melting Process Stability," J. Chem. Inf. Model., vol. 53, no. 9, pp. 1689–1699, 2013.
- [183] L. Yang, H. Gong, S. Dilip, and B. Stucker, "An Investigation of Thin Feature Generation in Direct Metal Laser Sintering Systems," in *Solid Freeform Fabrication*, 2014.
- [184] L. Parry, I. A. Ashcroft, and R. D. Wildman, "Understanding the effect of laser scan strategy on residual stress in selective laser melting through thermo-mechanical simulation," Addit. Manuf., vol. 12, pp. 1–15, 2016.
- [185] D. Zhang, W. Niu, X. Cao, and Z. Liu, "Effect of standard heat treatment on the microstructure and mechanical properties of selective laser melting manufactured Inconel 718 superalloy," *Mater. Sci. Eng. A*, vol. 644, pp. 32–40, 2015.
- [186] S. K. Mannan, E. L. Hibner, and B. C. Puckett, "Physical Metallurgy of Alloys 718, 925, 725, and 725HS for Service in Aggressive Corrosion Environments," *Corrosion*, p. 3125, 2003.
- [187] S. Safdar, A. J. Pinkerton, L. Li, M. A. Sheikh, and P. J. Withers, "An anisotropic enhanced thermal conductivity approach for modelling laser melt pools for Ni-base super alloys," *Appl. Math. Model.*, vol. 37, no. 3, pp. 1187–1195, 2013.
- [188] D. G. Thakur, B. Ramamoorthy, and L. Vijayaraghavan, "Effect of High Speed Cutting Parameters on the Surface Characteristics of Superalloy Inconel," in *World Congress on Engineering*, 2010, vol. III.
- [189] E. O. Olakanmi, R. F. Cochrane, and K. W. Dalgarno, "A review on selective laser sintering/melting (SLS/SLM) of aluminium alloy powders: Processing, microstructure, and properties," *Prog. Mater. Sci.*, vol. 74, pp. 401–477, 2015.
- [190] D. Rosenthal, "The Theory of Moving Sources of Heat and Its Application to Metal Treatments," *Trans. ASME*, 1946.
- [191] EOS GmbH, "Material Datasheet for EOS GP1 (SS 17-4)." 2009.

- [192] K. Devendranath Ramkumar et al., "Influence of filler metals and welding techniques on the structure-property relationships of Inconel 718 and AISI 316L dissimilar weldments," *Mater. Des.*, vol. 62, pp. 175–188, 2014.
- [193] K. B. Small, D. A. Englehart, and T. A. Christman, "Guide to Etching Specialty Alloys," *Adv. Mater. Process.*, pp. 32–37, 2008.
- [194] S. Narra, D. Christiansen, and J. Beuth, "Towards Enabling Spatial Control of Ti-6Al-4V Solidification Microstructure in the Electron Beam Melting Process," in *Solid Freeform Fabrication Symposium*, 2015.
- [195] C. Montgomery, J. Beuth, L. Sheridan, and Klingbeil, "Process Mapping of Inconel 625 in Laser Powder Bed Additive Manufacturing," in *Solid Freeform Fabrication Symposium*, 2015.
- [196] W. Tan and Y. C. Shin, "Analysis of multi-phase interaction and its effects on keyhole dynamics with a multi-physics numerical model," J. Phys. D. Appl. Phys., vol. 47, pp. 1–17, 2014.
- [197] B. A. Fisher, B. Lane, H. Yeung, J. Beuth, and S. Hall, "Toward Determining Melt Pool Quality Metrics via Coaxial Monitoring in Laser Powder Bed Fusion," *Soc. Manuf. Eng. Lett.*, 2018.
- [198] V. Gunenthiram, P. Peyre, M. Schneider, M. Dal, F. Coste, and R. Fabbro, "Analysis of laser-melt pool-powder bed interaction during the selective laser melting of a stainless steel," J. Laser Appl., vol. 29, no. 2, 2017.
- [199] L. E. Criales, Y. M. Arısoy, B. Lane, S. Moylan, A. Donmez, and T. Özel, "Laser powder bed fusion of nickel alloy 625: Experimental investigations of effects of process parameters on melt pool size and shape with spatter analysis," *Int. J. Mach. Tools Manuf.*, vol. 121, pp. 22–36, 2017.
- [200] U. Scipioni Bertoli, G. Guss, S. Wu, M. J. Matthews, and J. M. Schoenung, "In-situ characterization of laser-powder interaction and cooling rates through high-speed imaging of powder bed fusion additive manufacturing," *Mater. Des.*, vol. 135, pp. 385– 396, 2017.
- [201] M. Luo and Y. C. Shin, "Estimation of keyhole geometry and prediction of welding defects during laser welding based on a vision system and a radial basis function neural network," *Int. J. Adv. Manuf. Technol.*, vol. 81, pp. 263–276, 2015.
- [202] M. Khanzadeh, L. Bian, N. Shamsaei, and S. M. Thompson, "Porosity Detection of Laser Based Additive Manufacturing Using Melt Pool Morphology Clustering," in *Solid Freeform Fabrication*, 2016, pp. 1487–1494.
- [203] M. Khanzadeh, S. Chowdhury, M. A. Tschopp, H. R. Doude, M. Marufuzzaman, and L. Bian, "In-Situ Monitoring of Melt Pool Images for Porosity Prediction in Directed Energy Deposition Processes," *IISE Trans.*, Dec. 2017.
- [204] D. G. Lowe, "Object Recognition from Local Scale-Invariant Features," in *Proceedings of the International Conference on Computer Vision*, 1999.

- [205] National Aeronautics and Space Administration, "The Electromagnetic Spectrum." [Online]. Available: https://science.nasa.gov/ems. [Accessed: 10-Jun-2017].
- [206] TechImaging, "Photron FASTCAM Viewer," 2017. [Online]. Available: https://www.techimaging.com/downloads/. [Accessed: 18-Feb-2018].
- [207] MathWorks, "bwconncomp," 2017. [Online]. Available: https://www.mathworks.com/help/images/ref/bwconncomp.html. [Accessed: 22-Jun-2017].
- [208] S. Ly, A. M. Rubenchik, S. A. Khairallah, G. Guss, and M. J. Matthews, "Metal vapor microjet controls material redistribution in laser powder bed fusion additive manufacturing," *Nat. Sci. Reports*, vol. 7, pp. 1–12, 2017.
- [209] M. Jamshidinia, "Flaw Generation in Laser Powder Bed Fusion (L-PBF)," 2015. [Online]. Available: https://ewi.org/eto/wpcontent/uploads/2015/amcEWI716/03_Jamshidinia.pdf.
- [210] L. C. Wei, L. E. Ehrlich, M. J. Powell-Palm, C. Montgomery, J. Beuth, and J. A. Malen, "Thermal conductivity of metal powders for powder bed additive manufacturing," Addit. Manuf., vol. 21, no. December 2017, pp. 201–208, 2018.
- [211] R. Munroe, "XKCD," 2018. [Online]. Available: https://xkcd.com/. [Accessed: 17-Feb-2018].
- [212] B. L. DeCost, H. Jain, A. D. Rollett, and E. A. Holm, "Computer Vision and Machine Learning for Autonomous Characterization of AM Powder Feedstocks," J. Mater., vol. 69, no. 3, pp. 456–465, 2017.
- [213] D. G. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," Int. J. Comput. Vis., vol. 60, no. 2, pp. 91–110, 2004.
- [214] P. Dollar, "hogdraw." Piotr's Image & Video Toolbox version 2.41, Pasadena, CA, 2009.
- [215] V. Powell and L. Lehe, "Principle Component Analysis," 2015. [Online]. Available: http://setosa.io/ev/principal-component-analysis/. [Accessed: 18-Feb-2018].
- [216] N. Dalal and B. Triggs, "Histograms of Oriented Gradients for Human Detection," in *Conference on Computer Vision and Pattern Recognition*, 2005, pp. 886--893.
- [217] H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating local descriptors into a compact representation," in *IEEE Conference on Computer Vision & Pattern Recognition*, 2010, pp. 3304–3311.
- [218] MathWorks, "t-SNE," 2018. [Online]. Available: https://www.mathworks.com/help/stats/t-sne.html. [Accessed: 18-Feb-2018].
- [219] MathWorks, "fitcecoc," 2018. [Online]. Available: https://www.mathworks.com/help/stats/fitcecoc.html. [Accessed: 18-Feb-2018].
- [220] K. Bennett and C. Campbell, "Support Vector Machines: Hype or Hallelujah," *SIGKDD Explor.*, vol. 2, no. 2, pp. 1–13, 2000.

- [221] K.-B. Duan and S. S. Keerthi, "Which Is the Best Multiclass SVM Method? An Empirical Study," in *Multiple Classifier Systems*, 2005, pp. 278–285.
- [222] National Instruments, "Introduction to FPGA Technology: Top 5 Benefits," 2012. [Online]. Available: http://www.ni.com/white-paper/6984/en/. [Accessed: 06-Jul-2017].
- [223] S. Grantham, B. Lane, J. Neira, S. Mekhontsev, M. Vlasea, and L. Hanssen, "Optical design and initial results from NIST's AMMT/TEMPS facility," in *Laser 3D Manufacturing III*, 2016.
- [224] J. Mireles, C. Terrazas, S. M. Gaytan, D. A. Roberson, and R. B. Wicker, "Closed-loop automatic feedback control in electron beam melting," Int. J. Adv. Manuf. Technol., vol. 78, pp. 1193–1199, 2015.
- [225] W. M. Steen, Laser Material Processing, 3rd ed. London, U.K.: Springer-Verlag, 2003.
- [226] T. L. Bergman, A. S. Lavine, F. P. Incropera, and D. P. Dewitt, *Fundamentals of Heat Transfer*, 7th ed. John Wiley & Sons, 2011.
- [227] MathWorks, "normplot," 2017. [Online]. Available: https://www.mathworks.com/help/stats/normplot.html. [Accessed: 31-Dec-2017].
- [228] K. C. Mills, *Thermophysical Properties for Selected Commercial Alloys*. Woodhead Publishing, 2002.
- [229] A. M. Rubenchik, W. E. King, and S. Wu, "Scaling laws for the additive manufacturing," J. *Mater. Process. Technol.*, 2018.
- [230] Buehler, "Polishing Application Guide," 2015. [Online]. Available: https://www.buehler.com/Brochures/English/Grinding-Polishing/FN01465_0216_Polishing_Application_Guide_WEB.pdf. [Accessed: 29-Dec-2017].
- [231] EOS GmbH, "Compatibility DMLS Materials EOSINT M 290." 2016.
- [232] F. Sakuma and S. Hattori, "Establishing a Practical Temperature Standard by using a Narrow-Band Radiation Thermometer with a Silicon Detector," in *Temperature: its Measurement and Control in Science and Industry*, 2011, pp. 421–427.

Appendices

For clarification purposes, additional information and details have been provided in the appendices that follow. In the event of a discrepancy between the information or nomenclature in the appendices and the corresponding information in the body of the manuscript, the body text takes precedence.

Average AlSi10Mg 0LP melt pool cross-sectional measurements						
Sample Number	Beam Power (W)	Beam Velocity (mm/s)	Width (μm)	Depth (μm)	Area (mm²)	Aspect Ratio
1	100	200	71.6	18.4	0.00097	0.51
2	100	400	60.2	15.4	0.00065	0.51
3	100	600	58.6	14.6	0.00062	0.50
4	100	800	58.0	13.9	0.00056	0.48
5	100	1000	57.8	13.0	0.00056	0.45
6	200	200	141.0	60.0	0.00623	0.85
7	200	400	122.3	51.2	0.00478	0.84
8	200	600	115.6	46.1	0.00392	0.80
9	200	800	121.0	39.9	0.00366	0.66
10	200	1000	92.9	37.2	0.00263	0.80
11	200	1200	97.3	32.0	0.00239	0.66
12	200	1400	80.0	29.4	0.00181	0.74
13	300	400	367.2	294.6	0.07117	1.60
14	300	600	277.9	220.3	0.04234	1.59
15	300	800	237.6	171.4	0.02796	1.44
16	300	1000	187.9	115.4	0.01556	1.23
17	300	1200	142.4	61.6	0.00747	0.87
18	300	1400	101.7	45.6	0.00383	0.90
19	370	400	405.7	418.8	0.11406	2.06
20	370	600	326.2	282.7	0.06150	1.73
21	370	800	332.7	234.1	0.05431	1.41
22	370	1000	248.1	176.4	0.02953	1.42
23	370	1200	206.2	136.2	0.02039	1.32
24	370	1400	163.2	87.6	0.01036	1.07

Appendix A: Selected AlSi10Mg 0LP Micrographs and Measurements



AlSi10Mg 0LP Sample #1



AlSi10Mg 0LP Sample #2



AlSi10Mg 0LP Sample #3



AlSi10Mg 0LP Sample #4



AlSi10Mg 0LP Sample #5



AlSi10Mg 0LP Sample #6



AlSi10Mg OLP Sample #7



AlSi10Mg 0LP Sample #8


AlSi10Mg 0LP Sample #9



AlSi10Mg 0LP Sample #10



AlSi10Mg 0LP Sample #11



AlSi10Mg 0LP Sample #12



AlSi10Mg 0LP Sample #13



AlSi10Mg OLP Sample #14



AlSi10Mg 0LP Sample #15



AlSi10Mg 0LP Sample #16



AlSi10Mg 0LP Sample #17



AlSi10Mg 0LP Sample #18



AlSi10Mg 0LP Sample #19



AlSi10Mg 0LP Sample #20



AlSi10Mg 0LP Sample #21



AlSi10Mg 0LP Sample #22



AlSi10Mg 0LP Sample #23



AlSi10Mg 0LP Sample #24

Appendix B: Selected AlSi10Mg MLP Micrographs



AlSi10Mg MLP Sample #3

AlSi10Mg MLP Sample #4



AlSi10Mg MLP Sample #5



AlSi10Mg MLP Sample #7



AlSi10Mg MLP Sample #6



AlSi10Mg MLP Sample #8



AlSi10Mg MLP Sample #9



AlSi10Mg MLP Sample #11



AlSi10Mg MLP Sample #10



AlSi10Mg MLP Sample #12



AlSi10Mg MLP Sample #13



AlSi10Mg MLP Sample #14



AlSi10Mg MLP Sample #15



AlSi10Mg MLP Sample #16



AlSi10Mg MLP Sample #17



AlSi10Mg MLP Sample #19



AlSi10Mg MLP Sample #18



AlSi10Mg MLP Sample #20



AlSi10Mg MLP Sample #21



AlSi10Mg MLP Sample #23



AlSi10Mg MLP Sample #22



AlSi10Mg MLP Sample #24

Appendix C: Selected Cylindrical Sample Micrographs



Ti64 PS #1 Cylindrical Sample



Ti64 PS #2 Cylindrical Sample



Ti64 PS #3 Cylindrical Sample



Ti64 PS #4 Cylindrical Sample



In718 PS #5 Cylindrical Sample

Appendix D: Powder Bed Images Used for Testing



The above legend applies to all 20 testing powder bed images presented below.



 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$





 $\mathsf{X} \gets \mathsf{recoater} \ \mathsf{direction}$



 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$





 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$





 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$





 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$



 $\mathsf{X} \gets \text{recoater direction}$



Average In718 1LSB melt pool cross-sectional measurements						
Sample Number	Beam Power (W)	Beam Velocity (mm/s)	Width (μm)	Depth (µm)	Area (mm²)	Aspect Ratio
1	285	960	151.8	150.4	0.0149	2.0
2	100	200	177.4	166.4	0.0175	1.9
3	100	400	126.0	95.9	0.0079	1.5
4	100	600	103.2	69.1	0.0048	1.3
5	100	800	91.3	56.7	0.0035	1.2
6	100	1000	81.2	46.8	0.0025	1.2
7	150	200	227.1	293.5	0.0350	2.6
8	150	400	170.0	177.3	0.0185	2.1
9	150	600	147.3	109.2	0.0103	1.5
10	150	800	114.9	78.7	0.0062	1.4
11	150	1000	99.0	66.2	0.0047	1.3
12	150	1200	93.9	61.6	0.0039	1.3
13	200	200	265.4	403.5	0.0508	3.0
14	200	400	184.8	245.7	0.0260	2.7
15	200	600	167.1	152.5	0.0165	1.8
16	200	800	140.6	115.0	0.0106	1.6
17	200	1000	116.5	97.0	0.0078	1.7
18	200	1200	104.5	86.7	0.0061	1.7
19	250	200	277.7	521.4	0.0666	3.8
20	250	400	203.1	302.6	0.0336	3.0
21	250	600	174.9	199.2	0.0227	2.3
22	250	800	155.1	140.2	0.0143	1.8
23	250	1000	124.8	117.3	0.0105	1.9
24	250	1200	114.3	106.3	0.0085	1.9
25	250	1400	109.6	100.4	0.0076	1.8
26	300	400	228.2	282.7	0.0362	2.5
27	300	600	189.2	210.6	0.0256	2.2
28	300	800	169.2	167.5	0.0184	2.0
29	300	1000	149.2	140.8	0.0138	1.9
30	300	1200	122.5	130.9	0.0116	2.1
31	300	1400	115.3	116.4	0.0092	2.0
32	370	400	243.8	362.8	0.0483	3.0
33	370	800	171.8	227.8	0.0257	2.7
34	370	1000	158.9	182.8	0.0194	2.3
35	370	1200	130.9	157.8	0.0148	2.4
36	370	1400	123.0	142.9	0.0124	2.3

Appendix E: Selected In718 1LSB Micrographs and Measurements



In718 1LSB Sample #1



In178 1LSB Sample #2



In718 1LSB Sample #3



In718 1LSB Sample #4



In718 1LSB Sample #5



In718 1LSB Sample #6



In718 1LSB Sample #7



In718 1LSB Sample #8



In718 1LSB Sample #9



In718 1LSB Sample #10



In718 1LSB Sample #11



In718 1LSB Sample #12



In718 1LSB Sample #13





In718 1LSB Sample #15



In718 1LSB Sample #16



In718 1LSB Sample #17





In718 1LSB Sample #19



In718 1LSB Sample #20



In718 1LSB Sample #21



In718 1LSB Sample #22



In718 1LSB Sample #23



In718 1LSB Sample #24



In718 1LSB Sample #25





In718 1LSB Sample #27



In718 1LSB Sample #28



In718 1LSB Sample #29



In718 1LSB Sample #32



In718 1LSB Sample #33





In718 1LSB Sample #35



In718 1LSB Sample #36

Appendix F: Selected In-Situ High Speed Camera Images

Note that the relative spatial scale is preserved between all of the melt pool images presented below.



E2 melt pool PV #1

not available

E2 melt pool PV #5



E2 melt pool PV #9



E2 melt pool PV#13



E2 melt pool PV #17



E2 melt pool PV #2



E2 melt pool PV #6



E2 melt pool PV #10



E2 melt pool PV #14

not available

E2 melt pool PV #18



E2 melt pool PV #3



E2 melt pool PV #7



E2 melt pool PV #11



E2 melt pool PV #15



E2 melt pool PV #19



E2 melt pool PV #4



E2 melt pool PV #8

not available

E2 melt pool PV #12



E2 melt pool PV #16



E2 melt pool PV #20



E2 melt pool PV #21



E2 melt pool PV #25



E2 melt pool PV #22



E2 melt pool PV #26



E2 melt pool PV #30



E2 melt pool PV #34



E2 melt pool PV #23



E2 melt pool PV #24

not available

E2 melt pool PV #27



E2 melt pool PV #31



E2 melt pool PV #35

not available

E2 melt pool PV #28



E2 melt pool PV #32



E2 melt pool PV #36

not available

E2 melt pool PV #29



E2 melt pool PV #33